Bayesian Information Fusion for Precision Indoor Location

by

Andrew F. Cavanaugh

A Thesis
Submitted to the Faculty
of the
WORCESTER POLYTECHNIC INSTITUTE
in partial fulfillment of the requirements for the
Degree of Master of Science
in
Electrical and Computer Engineering
by

_____

2010

APPROVED:

_____

Professor David Cyganski, Major Advisor

_____

Professor R. James Duckworth

_____

Professor John A. Orr

Abstract

This thesis documents work which is part of the ongoing effort by the Worcester Polytechnic Institute (WPI) Precision Personnel Locator (PPL) project, to track and locate first responders in urban/indoor settings. Specifically, the project intends to produce a system which can accurately determine the floor that a person is on, as well as where on the floor that person is, with sub-meter accuracy. The system must be portable, rugged, fast to set up, and require no pre-installed infrastructure. Several recent advances have enabled us to get closer to meeting these goals: The development of Transactional Array Reconciliation Tomography(TART) algorithm, and corresponding locator hardware, as well as the integration of barometric sensors, and a new antenna deployment scheme. To fully utilize these new capabilities, a Bayesian Fusion algorithm has been designed.

The goal of this thesis is to present the necessary methods for incorporating diverse sources of information, in a constructive manner, to improve the performance of the PPL system. While the conceptual methods presented within are meant to be general, the experimental results will focus on the fusion of barometric height estimates and RF data. These information sources will be processed with our existing Singular Value Array Reconciliation Tomography (σART), and the new TART algorithm, using a Bayesian Fusion algorithm to more accurately estimate indoor locations.

Acknowledgements

To my family: My family has been right by my side assisting me in every possible way, and I cannot thank them enough.

To my sponsor: I would like to thank all of the entities without whom this work could never have been completed: Raytheon, the Department of Homeland Security, and the Federal Emergency Management Agency.

To my fellow team members: I owe a great deal of gratitude to Jorge Alejandro, Vincent Amendolare, Matthew Campbell, Matthew Lowe, Jamie Mitchell, Vasil Savov, and Benjamin Woodacre. Every single one of you has helped me to conduct experiments, write code, or given me the algorithms and techniques on which this thesis depends, in spite of your own busy schedules.

To my committee: I would also like to thank Professors Duckworth and Orr for being on my research committee, and for bringing me onto the PPL team and returning to an even more hands on role with the PPL project, respectively.

To my advisor: Finally, the person who helped me out countless times with theoretical questions, technical questions, practical questions, academic advising, paper writing, thesis writing, transportation, and being incredibly patient throughout all of it; I would like to thank Professor Cyganski.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis is written as part of the ongoing effort of the Precision Personnel Location (PPL) system being developed in the Electrical and Computer Engineering Department at the Worcester Polytechnic Institute (WPI). Our goal is to provide real-time accurate indoor location information for firefighters with a rapidly-deployable system, requiring no pre-installed infrastructure. In recent years the project has expanded to consider fire, police, and military personnel, as well as robots, in both urban and wilderness settings [10]. New algorithms [5] and hardware [9] have expanded our access to additional information correlated with, and potentially enhancing the estimation of the location of a first-responder. This thesis deals with incorporating this new information into the location solution, with specific experimental results to support the theory presented. Specifically, our goal is to develop a system that:

- Displays to an incident commander, the floor and 2D location of multiple responders in and around a building
- Always identifies the correct floor, with sub-meter accuracy in 2D location
- Requires *minimal setup time*
- Automatically configures itself with minimal user input
- Relays environmental and physiological information to a real time display
- Can operate over a 1Km^2 area

- Requires no pre-installed infrastructure

Throughout this document several acronyms are used. Each is defined in the text, but the acronyms are also provided in Table 1.1 as a quick reference. Similarly mathematical conventions are provided in Table 1.2 and variable names used in Table 1.3.

1.1 Precision Personnel Location Problem

On December 3, 1999, a fire broke out at the Worcester Cold Storage Warehouse, in Worcester, MA. This abandoned building was originally built for cold storage, which meant that it was very well insulated, and had few windows, all of which were boarded up at the time of the fire. Firefighters entered the building because it was thought that there were two individuals trapped inside the building at the time of the fire. This later turned out to be false. The warehouse complex consisted of a six story cold storage facility and a second building that was connected to it. Both structures consisted of six stories, plus a basement, but it is unclear if the layout of these structures was known at the time of the fire [23, 6]. The magnitude of the fire led to many problems with the organization of radio communications, which made keeping track of every team of firefighters impossible. The situation was further complicated by the fact that some firefighters had entered from the ground and others had entered through the roof, giving them no common floor reference. Within two hours of sounding the first alarm, six firefighters had lost their lives because they were unable to find a safe exit from the building.

This tragedy revealed the need for a way of locating firefighters inside of a building. The Worcester Cold Storage fire occurred in a building that was abandoned, and therefore not compliant with fire codes; but it was known to be inhabited, requiring rescue teams to be sent into the building. The building layout presented a significant challenge for search and rescue personnel.

While our system is being developed with firefighters as our primary users, police, corrections, and military personnel could also benefit from location technology; there are numerous applications from search and rescue scenarios, to friendly fire avoidance. In all cases the system needs to be deployable on site, requiring no pre-existing infrastructure. The current system has been successfully demonstrated in tracking autonomous vehicles in large(1 Km^2) outdoor settings.

1D	One-Dimensional
2D	Two-Dimensional
3D	Three-Dimensional
ADC	Analog to Digital Converter
DAC	Digital to Analog Converter
DFT	Discrete Fourier Transform
DR	Dead Reckoning
DSS	Direct State Space
DSP	Digital Signal Processor
DWG	Digital Waveform Generator
FFT	Fast Fourier Transform
FPGA	Field Programmable Gated Array
GAC	Geometric Autoconfiguration
GDOP	Geometric Dilution of Precision
GPIO	General Purpose Input/Output
GPS	Global Positioning System
HVAC	Heating Ventilation Air Conditioning
IC	Incident Commander
IMU	Inertial Measurement Unit
I/Q	In-phase Quadrature-phase
LAN	Local Area Network
LO	Local Oscillator
LOS	Line of Sight
MCWB	Multicarrier Wide Band
MSPS	Mega Samples Per Second
NIMU	Navigation Inertial Measurement Unit
NLOS	Near Line of Sight
PCB	Printed Circuit Board
PDF	Probability Density Function
PPL	Precision Personnel Location
RF	Radio Frequency
RMS	Root Mean Square
RSSI	Received Signal Strength Indicator
RX	Receiver
σART	Singular Value Array Reconciliation Tomography
SNR	Signal to Noise Ratio
SRM	Short Range Radio Module
SVD	Singular Value Decomposition
TART	Transactional Value Array Reconciliation Tomography
TCVR	Transceiver
TDM	Time Division Multiplexing
TDOA	Time Difference of Arrival
TOA	Time of Arrival
TX	Transmitter

Table 1.1: Acronyms

Time and frequency	$x(t) \overset{\mathcal{F}}{\Leftrightarrow} X(\omega) = X(2\pi f)$
Imaginary unit	$j = \sqrt{-1}$
Unit-Impulse Function [17]	$\delta(t) = 0 \quad \text{for } t \neq 0$ $\delta(t) \quad \text{undefined for } t = 0$ $\int_{t_1}^{t_2} \delta(t) = \left\{ \begin{array}{ll} 1, & \text{if } t_1 < 0 < t_2 \\ 0, & \text{otherwise} \end{array} \right\}$
Singleton	x
Vector	$\mathbf{x}$
Matrix	$\mathbf{X}$
Transpose	$\mathbf{X}'$
Conjugating Transpose	$\mathbf{X}^H$
Approximation of x	$\tilde{x}$
Entrywise, Hadamard or Schur product	$\mathbf{Z} = \mathbf{X} \circ \mathbf{Y}$, $(\mathbf{Z})_{a,b} = (\mathbf{X})_{a,b}(\mathbf{Y})_{a,b}$

Table 1.2: Mathematical Conventions

RF-Based Solution In response to this problem, several faculty, in the Electrical and Computer Engineering department, at the Worcester Polytechnic Institute, conceived of an RF-Based indoor positioning system: the Precision Personnel Locator (PPL). At that time, the available technology made RF-based systems the most attractive option. Silicon inertial measurement units (IMUs) were still in their infancy. Furthermore, IMU and dead reckoning (DR) techniques, which are employed by inertial navigation systems, suffer from the accumulation of errors that result from temperature sensitivity, $\frac{1}{f}$ noise [21], and imperfections in the manufacturing process that can be difficult to calibrate away. These contributions are exacerbated by the fact that the location solution is obtained by integrating the acceleration measurements in order to get velocity estimates while the velocity estimates are integrated to obtain position estimates. This double integration means that any bias error, ϵ, in an accelerometer measurement, will cause a position error that will grow over time as $\frac{1}{2}\epsilon t^2$. While RF solutions are also prone to errors, the errors do not accumulate, and in the absence of noise, do not vary with time at a given point in space.

The initial PPL system concept made use of a multicarrier signal structure, and new super resolution techniques from the field of radar to estimate TOA/TDOA. A TDOA approach was initially decided upon, since synchronization was only needed with the stationary units, not the mobile locators. The traditional algorithms for estimating TDOA [7] were evaluated with the RF system. While

Speed of light in a vacuum [meters/second]	c_0
σART data	D_σ
TART data	D_T
Barometric data	D_B
Prior information	I
Time [seconds]	t
Discrete time [samples]	k
Number of time samples	K
Frequency [Hz]	f
Angular frequency [radians/second]	ω
Signal period [seconds]	T
Frequency spacing between carriers	Δf, $\Delta\omega$
Signal index	n
Number of signals	N
Carrier index	m
Number of carriers	M
Reference unit index	p
Number of reference units	P
Mobile unit index	q
Number of mobile units	Q
Time delay for signal n	t_n
Time delay for signal n and Reference Antenna p	$t_{n,p}$
Attenuation for signal n	γ_n
Attenuation for signal n and Reference Antenna p	$\gamma_{n,p}$
Vector of carrier frequencies	$\mathbf{f}$ or $\boldsymbol{\omega}$
Centered vector of carrier frequencies	$\dot{\boldsymbol{\omega}}$
Transmitted signal	$x(t)$,$X(\omega)$, $x(k)$, $\mathbf{x}$
Received signal	$r(t)$,$R(\omega)$,$\mathbf{r}$
Channel response	$h(t)$,$H(\omega)$,$\mathbf{h}$
Sample clock frequency	f_s
RF mixer frequency	f_c
Random time-offset	$\tilde{\tau}$
Random phase offset	$\tilde{\theta}$
Position vector	$\mathbf{v}$
Mobile unit position	$\mathbf{v}_l$
Mobile unit position test value	$\overset{\#}{\mathbf{v}}_l$
Mobile unit position estimate	$\overset{\star}{\mathbf{v}}_l$
Reference unit position	$\mathbf{v}_p$
Received signals matrix	$\mathbf{R}$
Rephased received signals matrix	$\mathbf{R}'$
Ideal synchronization rephased received signals matrix	$\mathbf{R}'_{ideal}$
Received signals matrix rephased to correct location	$\overset{\star}{\mathbf{R}}'$
Time delay test value for Reference Antenna p	$\overset{\#}{t}_{0,p}$
Rectified received signal	$\mathbf{s}$
Ideal TART metric function	$f(\mathbf{v})$

Table 1.3: Variable Names

initial proof-of-concept tests with both ultrasound, and RF, in environments offering direct paths for the ranging signals, proved very successful, it soon became apparent that the problem of multipath was not going to be solved using traditional methods. A novel algorithm called σART [15] was developed to fuse all available RF signal information, rather than simply fusing computed ranges from individual antennas. This algorithm proved successful in mitigating the effects of disruptive levels of multipath. The TART [5] algorithm is the newest algorithm, and, like σART , it considers all received data when computing a solution. TART was developed to further extend the ability to mitigate multipath problems in even higher multipath channels. Unlike σART, TART requires TOA-like synchronization, which has led to the development of new locator hardware that can both transmit and receive [9], to support synchronization.

1.2 Rapid Deployment

In order for the PPL system to be utilized by first responders, on location, the system cannot require substantial additional time or effort, by the firefighters, to deploy. There can also be no requirement for pre-installed infrastructure, since an incident can happen anywhere at any time. To move towards this goal, the PPL system has been reconfigured to be very easy to deploy. Antennas that are pre-hung on ladders, or similar structures which can be quickly thrown up against, or next to a building have been the subject of recent research. The system can then be automatically configured by the incident commanders computer, using the Geometric Autoconfiguration (GAC) algorithm, which uses the receiving antennas to locate one another, on a relative coordinate system, with signal exchanges [25]. Although this setup improves the usability of the system, recent tests have shown a degradation of our location performance with this configuration, largely caused by the Geometric Dilution of Precision (GDOP) that results from deploying the antennas in a nearly planar configuration.

This thesis will present several techniques that will ameliorate this degraded performance. Chapter 2 summarizes the work done on the PPL system to date, as well as providing justification for assumptions that will be used to develop the theory of our Bayesian Fusion Approach, which is developed in Chapter 3. Chapter 4 explains specific advances in our system hardware and algorithms that make the approaches in this thesis possible to implement. Simulated results of the Bayesian

Fusion Algorithm are discussed in Chapter 5. The results of the implementation of the Bayesian Fusion algorithm are shown in Chapter 6. The impact of this work on the overall performance of the system will be considered, as well as future work that can be done along these lines in Chapter 7.

Chapter 2

WPI Precision Personnel Location System

This chapter will discuss the WPI Precision Personnel Locator (PPL) hardware and algorithms. The current, revision 2, hardware has capabilities suited for rapid deployment of our system, that were not present in the classical, revision 1, hardware. The two most current algorithms are also discussed to provide the required background for understanding the strengths of our Bayesian Fusion Algorithm. Singular Value Array Reconciliation Tomography, or σART , has been the algorithm employed by the PPL system since 2006. The newer TART (Transactional Array Reconciliation Tomography) algorithm builds on σART to achieve a higher level of synchronization, as well as a more computationally efficient location metric. The hardware required to use the TART algorithm (revision 3) is still under development, but preliminary tests were conducted with our revision 2 hardware, as explained in section 2.3.

2.1 System Hardware

2.1.1 Classical Configuration

Revision 1 Hardware The classical revision 1 WPI PPL system, as shown in Figure 2.2, consists of a mobile locator that transmits a multicarrier signal. The signal is picked up by antennas that are connected to receivers on the outside of a building. These receivers down mix the signal and send

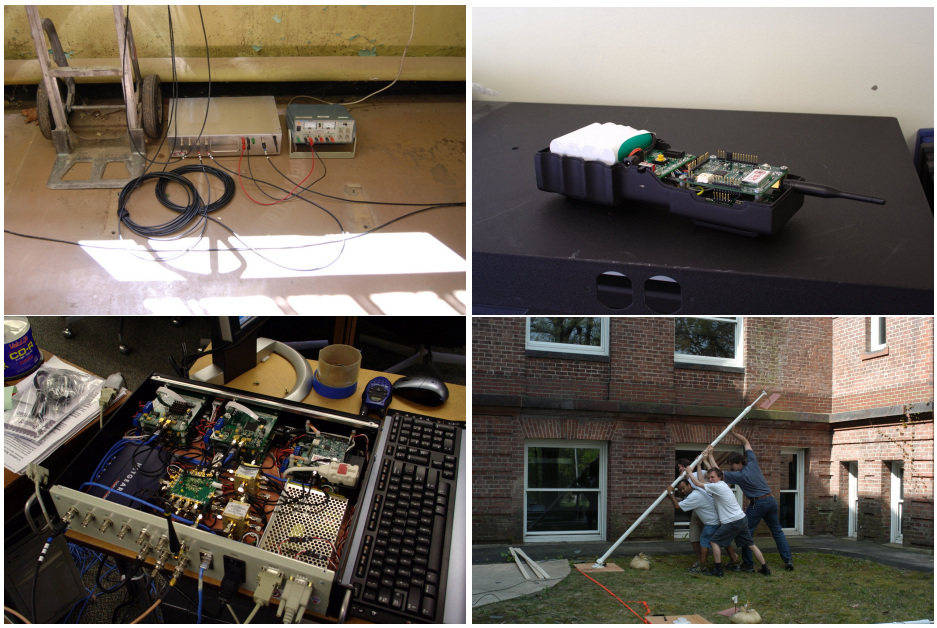

Figure 2.1: Classical System: Baseband cables running between receiver boxes, classical locator hardware, classical ADC stack, difficult setup

a baseband signal to a box which contains 5 ADCs that is connected to a base station computer. The ADCs are all connected to the same clock, which ensures synchronization of the sample clocks within one clock period. To capture data from 16 receiving antennas, four of these ADCs were used in a time-division multiplexing (TDM) configuration. The fifth ADC was connected to a reference antenna, that actively captures signal samples during all of the TDM divisions, which allows us to track the sample clock drift across TDM divisions. Once the drift is removed, the signals are fused and processed, on the base station computer, with the σART algorithm running in either real time[12] or post processed in a Matlab environment. This system required baseband cables to be run between the receive arrays and the central ADC stack. Having all of the ADCs on a single clock was advantageous, but a realistic scenario would never accommodate the setup time, or the danger of running these cables around a structure.

Revision 2 Hardware Our revision 2 hardware, as shown in Figure 2.3, allows us to operate in the classical configuration without the need for cabled links, traditionally needed for communications and synchronization within the PPL system, between transceiver boxes. The transceiver boxes are operated primarily as receivers, transmitting only when we need to renew our coarse synchroniza-

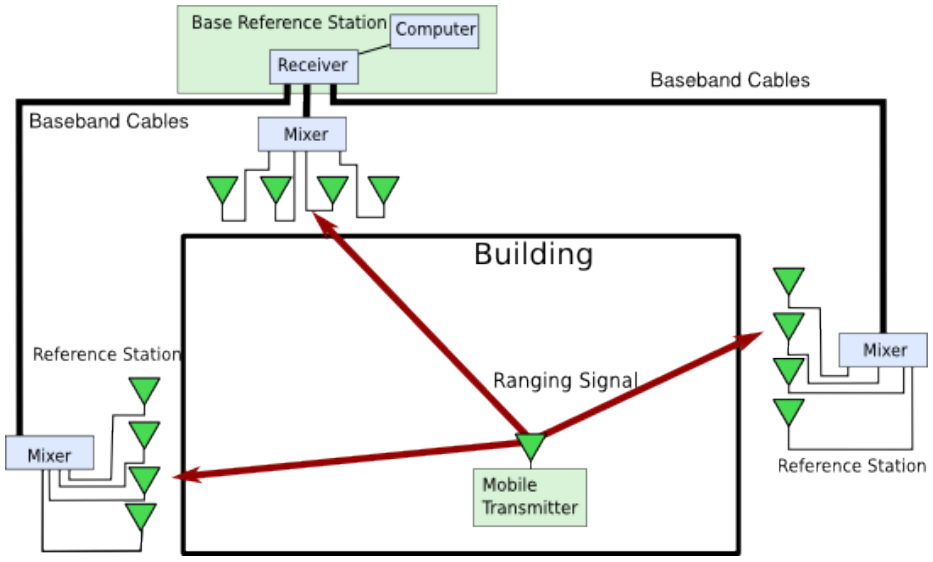

Figure 2.2: σART Revision 1 System Layout [5]

tion[1].

2.2 Singular Value Array Reconciliation Tomography (σART)

2.2.1 σART Theory

Currently, the PPL system utilizes the σART algorithm, a unique algorithm that was developed at WPI [15]. This algorithm is a TDOA-like RF based approach, which considers data from all of the receiving antennas as one set, rather than performing individual 1D ranging estimates. The σART algorithm is performed on the received signal sample vector, from a set of p receiving antennas. The received signal, is the multicarrier wide band (MCWB) signal, transmitted by the locator, X, convolved with the channel response between the locator, and a receiving antenna, p. σART is performed on a matrix, $\mathbf{R} \in \mathbb{C}_{N\times p}$, whose columns are the N-point FFT of the received signals from the p receive antennas. Each column of $\mathbf{R}$ is, therefore, equal to $\mathbf{r}_p(\omega) = X(\omega)H_p(\omega)$. The MCWB signal structure is a sum of unmodulated sinusoids [5], evenly spaced in frequency, as shown in Equation (2.1). Here, we are ignoring initial phase, as well as any time offsets introduced by the

[1]Our atomic clocks allow us to perform coarse synchronization once per deployment, traditionally we have needed to coarse sync once every 5 minutes.

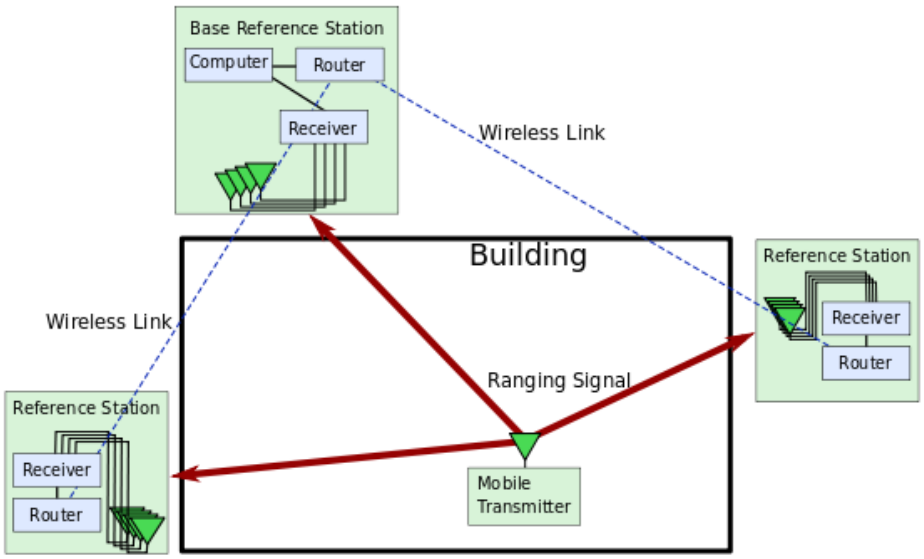

Figure 2.3: σART Revision 2 System Layout [5]

hardware.

$$X(\omega) = \sum_{n=0}^{m-1} \delta(\omega - (\omega_0 + n\Delta\omega)) \tag{2.1}$$

If we consider phase offset in our mixers, $\tilde{\theta}(t)$, and sample clock drift, $\tilde{\tau}(t)$, we obtain a more accurate expression for our transmitted signal in (2.2). Here, we are assuming that the time-dependency of these offsets is slow enough to represent them as constants for short windows of time.

$$X(\omega)' = \sum_{n=0}^{m-1} \delta(\omega - (\omega_0 + n\Delta\omega))e^{-j\omega\tilde{\tau}-\tilde{\theta}} \tag{2.2}$$

The received signal is therefore:

$$R_p(\omega) = X(\omega)H_p(\omega)e^{-j\omega\tilde{\tau}_p-\tilde{\theta}_p} \tag{2.3}$$

2.2.2 σART Metric

In order to estimate a transmitter location from the received signal matrix, a metric is evaluated at every point in the "scan grid", a discrete set of points in the solution space. The point where the metric is maximized is the σART algorithm's estimate of the locator's position. To evaluate the σART metric at a given point, the first step is to 're-phase' the received signal matrix:

$$\mathbf{R}' = \left[\mathbf{r}_1 \circ e^{j\omega\overset{\#}{t}_{0,1}} \dots \mathbf{r}_p \circ e^{j\omega\overset{\#}{t}_{0,p}} \dots \mathbf{r}_P \circ e^{j\omega\overset{\#}{t}_{0,P}}\right] \tag{2.4}$$

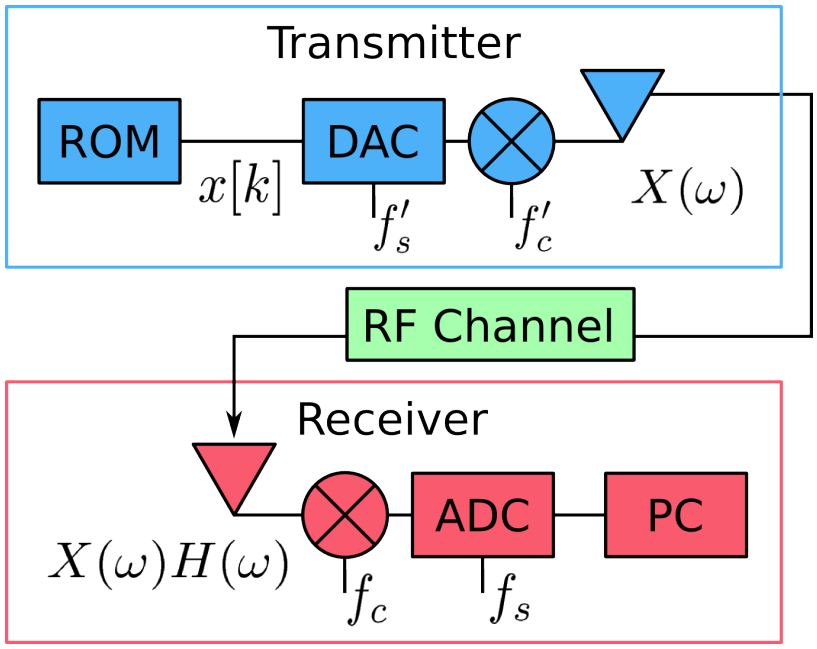

Figure 2.4: Block Diagram of Tx and Rx [5]

Where $\overset{\#}{t}_{0,p}$ is the ideal time offset between the scan grid location and the receiving antenna, p. This process is described in detail in [1, 5, 25]. The channel's impulse response will take the form of a set of delayed and scaled impulse functions. The transfer function, $H(\omega)$, is shown in (2.5), for the case of N_p paths.

$$H_p(\omega) = \sum_{n=1}^{N_p} \gamma_{n,p} e^{jwt_{n,p}} = \gamma_{1,p} e^{jwt_{1,p}} + \sum_{n=2}^{N_p} \gamma_{n,p} e^{jwt_{n,p}} \tag{2.5}$$

In general, the multipath components are weakly correlated, allowing the σART algorithm to largely ignore their contribution. At the correct scan location, the re-phased direct path components will be highly correlated, and the matrix, $\mathbf{R}'$ will become singular. It is shown in [5] that the energy of the received data matrix is not altered by re-phasing. The non-ideal phase offsets from (2.3) are also shown to be ignored by the σART algorithm. Since energy of the $\mathbf{R}$ matrix is conserved at every point in the scangrid, we can use the first singular value of the re-phased matrix as the σART metric.

2.2.3 σART metric as a PDF

To utilize the σART algorithm in our sensor fusion, we must be able to treat the information provided by it in a probabilistic manner. Simulations and perturbation analysis have shown a direct

correlation between the relative values of the σART metric, and relative frequency distribution of the location solution with noisy data [5].

Figure 2.5 shows a σART metric and a histogram of the simulated position estimates from 1000 trials of the same σART simulation with a -6dB SNR [5]. The metric functions both have a peak at the same point. The peaks are also elongated in the same direction. The relative frequency histogram lacks detail in the lower likelihood parts of the image because it was conducted for a finite number of trials, and these areas correspond to very low likelihoods.

In order to condition the metric for application as a probabilistic measure within the context of estimation theory, we normalize the final metric function, such that the metric values $\in [0, 1]$. To achieve this, we subtract off the minimum value and then divide by the maximum value. The metric is now a normalized discrete likelihood function for the location of a first responder. Chapter 3 describes how the information from these likelihood functions can be formed into a single likelihood function.

2.3 Transactional Array Reconciliation Tomography (TART)

2.3.1 TART Theory

The TART algorithm is the latest addition to our RF location and tracking system. With this new transactional approach, we are able to achieve TOA-like synchronization. This means that there is no unknown constant time offset across the columns of the received data matrix, Equation (2.3) now in effect becomes Equation (2.3)

$$R_p(\omega) = X(\omega)H_p(\omega)e^{j\tilde{\theta}_p} \tag{2.6}$$

The means to accomplish this synchronization is described in detail in [5]. New hardware is currently being developed to allow us to use TART in real time [9], but post-processing has enabled us to perform TART location with our current hardware [5], using the transceiver boxes as both the mobile, and the reference radios.

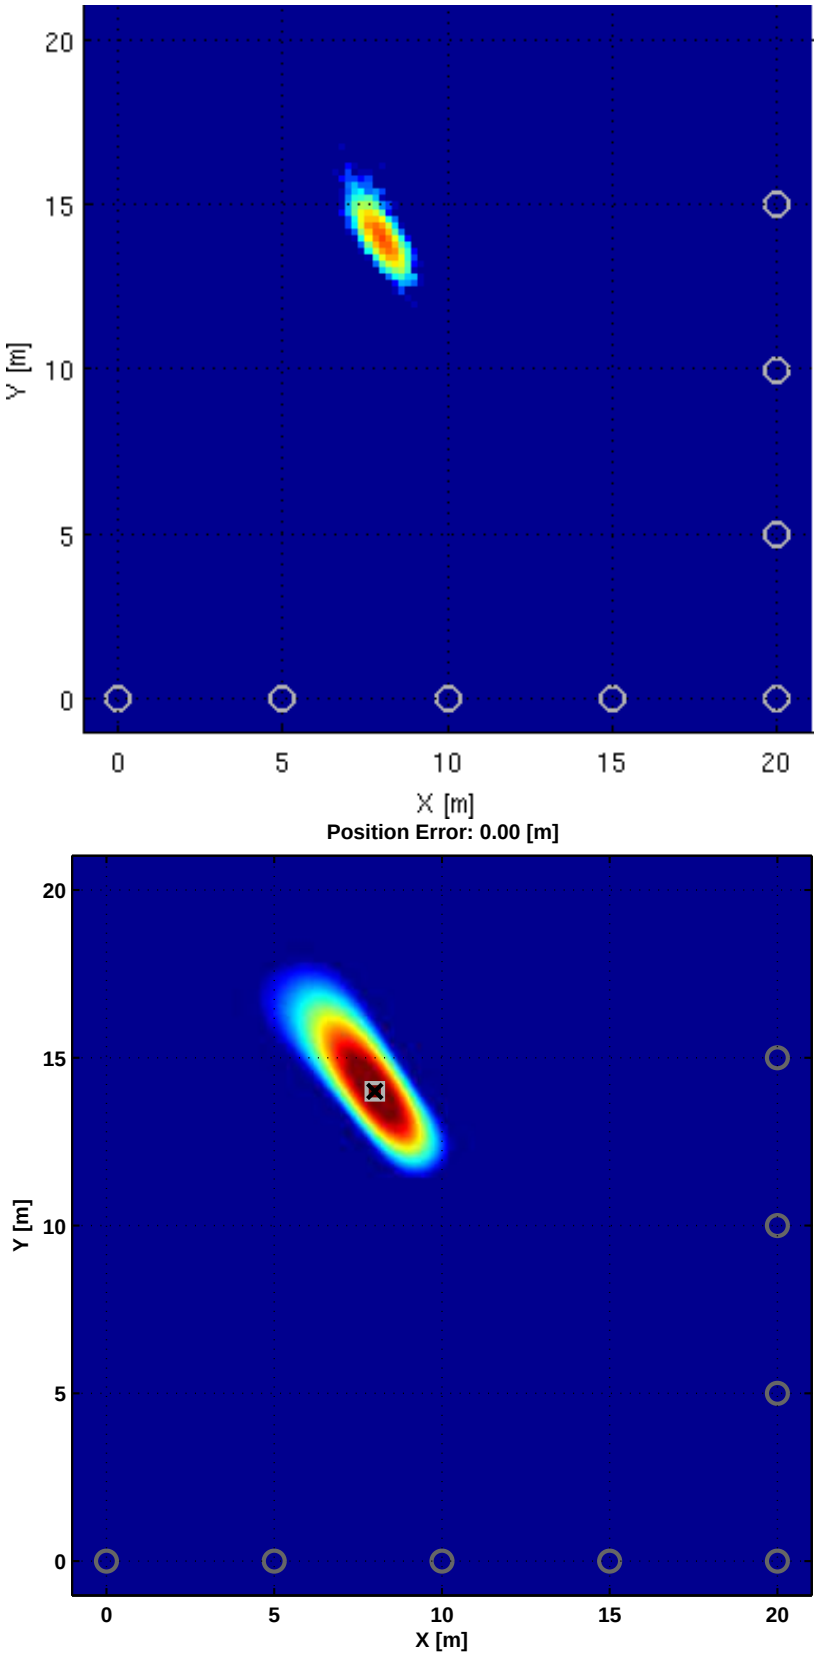

Figure 2.5: σART solution histogram(top) for 1000 Monte Carlo trials with noisy data [5] and σART metric function with highest values highlighted(bottom)

2.3.2 TART Metric

Having successfully removed the random time offsets from the received data matrix, the re-phasing process is carried out as detailed in Equation (2.4). At this point, the metric can be evaluated at every point in the scan grid. While σART required calculation of the first singular value of the re-phased matrix, TART only requires summation across the rows of the re-phased matrix. Since the 'DC' components will add constructively, and the true scan grid location will have the greatest 'DC' contribution. The computation time required for evaluating the TART metric is more than 10 times less than that required for σART [5].

2.3.3 TART metric as a PDF

The TART metric can also be turned into a normalized likelihood function defined on the scan grid. The same analysis from [5] shows that the relative frequency statistics are proportional to the metric magnitudes. We can thus normalize the metric so that the values are $\in [0, 1]$, allowing us to employ the same processing techniques that we would with the σART metric. Scaling the metric preserves the proportionality relationship.

Figure 2.6 shows a TART metric function and a histogram of the simulated position estimates from 1000 trials of the same TART simulation with a -6dB SNR [5]. Once again, the areas of high likelihood are proportional to the relative frequency histogram. The low-likelihood areas were not sufficiently likely to have been populated by events in great numbers over the 1000 trials computed.

2.4 Current Hardware

The sensor fusion algorithm presented in this thesis is made possible by several new additions to the PPL system. PPL hardware revision 2 came on line in later 2008. The most notable difference between the original hardware and our current hardware is the addition of on-board ADCs and FPGA based processors to our receiver boxes, as well as a DAC and RF front-end, which gives us full transceiver[2] capabilities. The new synchronization schemes, presented in a previous thesis [1], allow our transceivers to operate without any cables running between boxes for clock distribution,

[2]When referring to the device, receiver and transceiver will be used interchangeably.

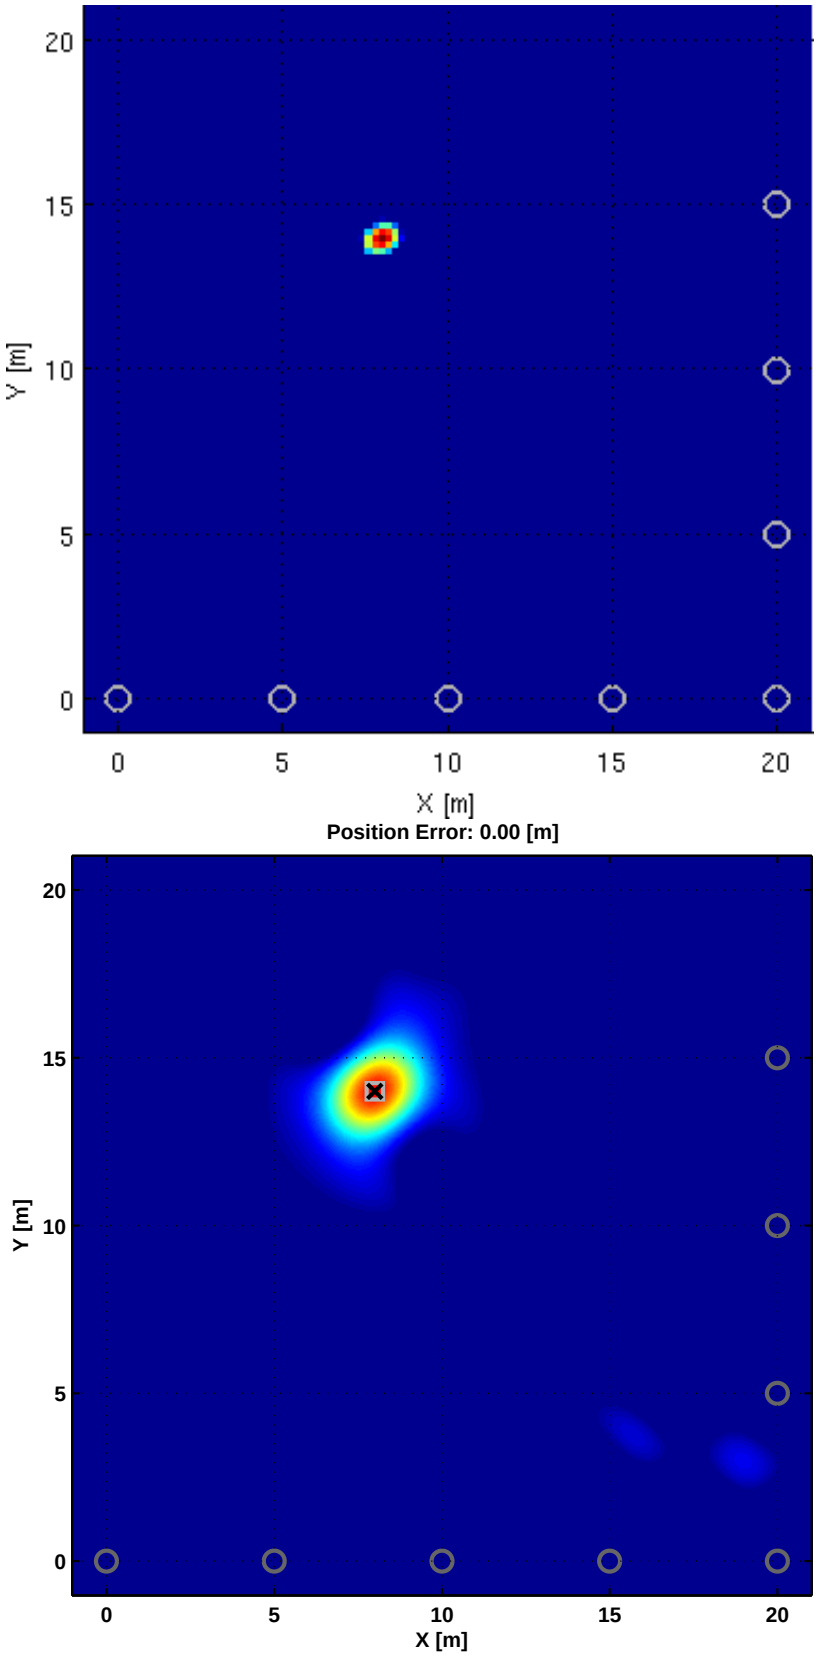

Figure 2.6: TART solution histogram(top) for 1000 Monte Carlo trials with noisy data [5] and TART metric function with highest values highlighted(bottom)

or base-band signal transfer. In the past year, Rubidium atomic clocks have been added to the transceivers to reduce the number of coarse synchronizations that must be performed. These locators (transmitters) were also outfitted with pressure, temperature, and inertial sensors, as well as a high power RF amplifiers. We also now have greater flexibility in the implementation of our system, since both the transceiver boxes, as well as the locators, are controlled by soft microprocessors implemented on FPGAs.

The PPL transceiver is a rapidly configurable RF transceiver and DSP. This is made possible by the Xilinx Vertex IV LX100 FPGA, which is designed for high speed signal processing applications. The FPGA is the interface between the RF components and the communications links between other boxes. The PPL transceiver has a four channel RF front-end with independent 14-bit, 400 Mega-sample per second (MSPS), ADCs[3] that allows for the simultaneous capture of data from up to four antennas that can be connected to a single transceiver box. The single 16-bit, 400 MSPS, DAC[4] and transmitting front end can transmit from any one antenna at a time. The communications interfaces include a 915 MHz Xemics XE1205 RF module, used for communications between transceivers and locators, as well as an Ethernet PHY, which is currently used to send data between transceiver boxes. Routers, on a local area network (LAN) with the base station computer, pick up the wireless data from the remote transceiver boxes, and forward them to the computer for processing. Although the individual transceiver boxes are capable of high speed signal processing, we forward the raw ADC data to the base station so that we can use data captured during experiments in use-case settings for laboratory simulations. The interior of the transceiver box is shown in Figure 2.7.

The PPL locator consists of two PCBs in a housing that is designed to be mounted on a person or vehicle. Each side of a board has a specific function. The four modules are: a Digital Waveform Generator (DWG), a 915 MHz. data-channel, the PPL Short Range Radio Module (SRM), and a power board. The Digital Waveform generator is a transmitting RF front end, driven, through a DAC, by a Xilinx Spartan III FPGA. The DWG can dynamically generate our PPL multicarrier signal. An on-board ROM stores up to four different signal configurations, or waveforms. The data-channel, which is identical to the data-channel in the transceiver boxes, is also controlled by

[3]Texas Instruments ADS5474
[4]Analog Devices AD9726

Figure 2.7: PPL Transceiver Box

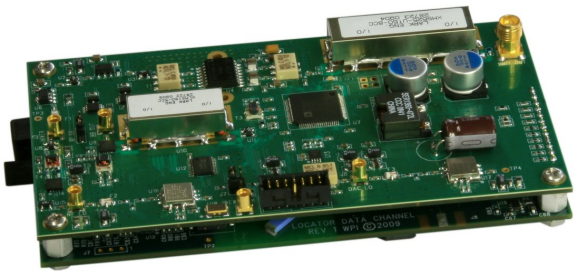

Figure 2.8: PPL Locator Boards: From Top: DWG, Data-Channel, PPL SRM, power

a Spartan III FPGA. The data-channel board also has additional sensors, including: temperature sensor, barometric sensor, 3-axis accelerometer, as well as interfaces for an NIMU, and general purpose input/output (GPIO) headers. The PPL SRM is a short-range wireless communications interface for the physiological monitoring equipment that is worn by a firefighter, or other first responder. This board uses a 915MHz. TI CC1101 ultra low power RF transceiver module, with an on-chip antenna. The SRM is controlled by the FPGA as well. Finally, the power board controls the battery charging, and contains switching regulators to achieve the required voltages used in the FPGA and RF hardware. Figure 2.8 shows the locator boards.

2.4.1 Synchronization Considerations

The fundamental concept that all RF based positioning systems rely on is the relationship between distance and time. RF signals in free space propagate at the speed of light, $c_0 = 299792458\frac{m}{s}$, so $distance = c_0 \cdot time$. Since the speed of light is very large compared with the distances that we would like to measure, we need to have extremely accurate clocks. A useful approximation to visualize this concept is that a 1 nanosecond error in time of arrival is equivalent to one foot of error in range, which means that we need synchronization to within approximately ±0.5ns in order to obtain positioning errors approximately under 1 foot under ideal RF conditions. There are also logistical synchronization considerations. In order for hardware to communicate effectively, and execute the correct commands on time they need to have a common time basis. For a TDOA system, only the receivers need to be synchronized with one another, to get TOA information, both the reference and mobile hardware need to be on a common time basis.

2.4.2 Coarse Synchronization

Coarse synchronization is needed to synchronize the ADC sample clocks to within one frame length[5]. A frame, in our context is one complete period of our 2048 sample period digital time domain signal from the DAC. Without coarse sync the transceiver boxes might possibly capture their data from different periods of the transmitted signals. If this were to happen then the data captured at each box would appear synchronized after “fine synchronization”, but the received data would actually be off by an integer number of periods. While the position of the locator would not

[5] 20.48 μs

change much on the order of microseconds, the channel response may be changing rapidly from the presence of noise and interference. If the frame count error were large enough, the received frames might even have been transmitted from different locators if multiple locators were in a TDM mode. The coarse sync. procedure is illustrated in Figure 2.9. The boxes maintain coarse synchronization through one of two methods. The traditional σART system relies on the stability of rubidium atomic clocks to maintain synchronization between coarse syncing events, which takes close to a minute to complete with the current hardware. The new TART system, however, can work with crystal oscillators, because transactional synchronization can be performed on the fly.

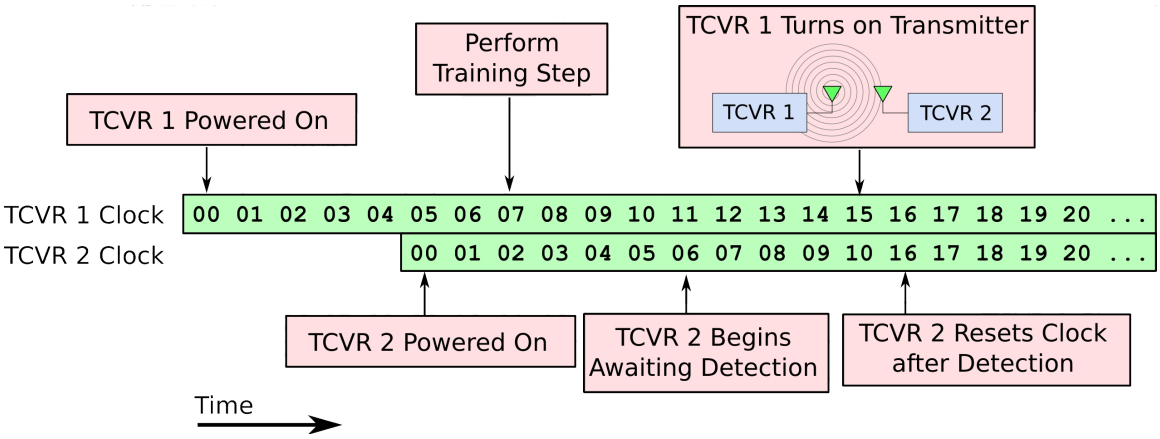

Figure 2.9: Coarse Synchronization Diagram [3]

Figures 2.10 through 2.12 shows the different ways that we can keep our hardware synchronized. In the first hardware revision, there was no need to perform coarse sync because all of the ADCs were co-located and on a common clock. In the Revision 2 system, where we had to coarse sync once per day of testing, we co-located the transceivers, sent a reset signal to all of the boxes, and their rubidium atomic clocks drifted slowly enough that this coarse synchronization would not need to be repeated for that test. The transceivers could then be moved about freely, as long as they were not powered off. The Revision 3 TART locator is designed to be able to perform an RF coarse sync wirelessly. We experimented with RF based coarse sync in Revision 2 tests, but found that it was easier to sync once and use the stability of the atomic clocks to maintain synchronization. The TART locators will allow us to rapidly perform the RF coarse sync procedure, and constantly correct for our clock drift. The specifics of this new synchronization are explained in references [5, 9]. It is important to note that the ability to maintain synchronization in the Rev 3 hardware is the result of

better processing power and algorithms, with lower quality clocks than the Rev 2 system.

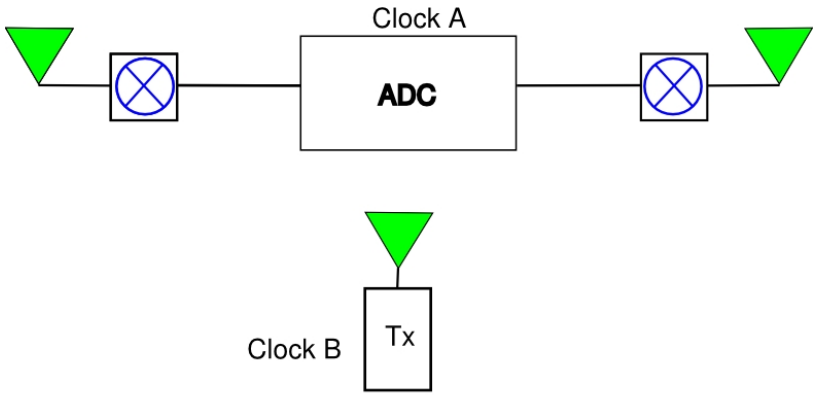

Figure 2.10: Rev 1 system had all ADCs on a common clock, requiring no coarse synchronization

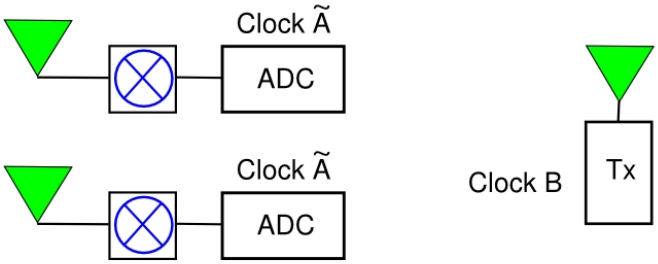

Figure 2.11: Rev 2 system coarse synced to a common reference, Clock A, $\tilde{A}$ represents the slowly drifting clocks

We also need coarse synchronization to track multiple responders, where we use time division multiplexing (TDM), so that only one locator is transmitting at a given time. The receivers collect the data from the locators during the prescribed time slots, and the σART algorithm, executing on the base station computer, is used to determine the position of each locator, which can then be displayed with an ID determined by the time slot corresponding to that locator.

2.4.3 Fine Synchronization

In the σART algorithm, after coarse synchronization is complete, the remaining sample clock time offsets are removed in a process that we call fine synchronization. Since the transceivers are

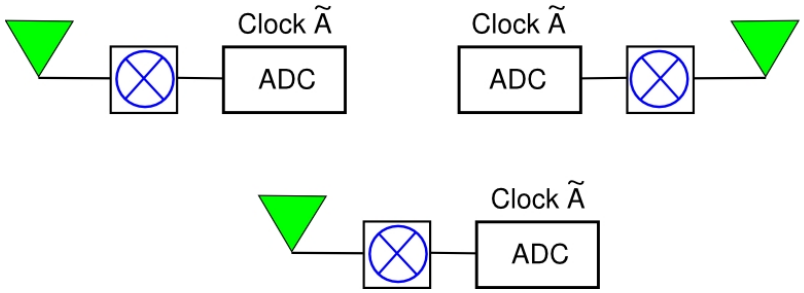

Figure 2.12: In the proposed Rev 3 system, transactional coarse sync allows crystal clocks to hold $\tilde{A}$ over short time periods

only synchronized on the order of $\pm 20.48 \mu s$, we must solve for more exact time offsets between the transceiver clocks, in order to get accurate location solutions[6]. This is done by initializing the locator at a known point, performing a σART capture, and solving for the time offsets that produce the maximal σART metric at that point. Clock drift is another problem, but the use of a stationary locator or transceiver allows us to cancel out time-varying effects, as detailed in [2]. The bidirectional TART approach takes fine synchronization into account. Since the transaction occurs in a very short period of time, the clocks and channel can be assumed constant, and the mobile and reference nodes can be assumed stationary with respect to one another. We can use the data from the transaction to solve for the time offset between the transmitted signal and the received signal. Since the signal path is the same in both directions, and the clocks have not drifted significantly, we are able to extract the clock offset. Much more detail on this topic, as well as a novel approach to clock drift tracking can be found in [5].

2.4.4 Hardware Under Development

To perform TART in real time, new locator hardware is being developed, with full transceiver capabilities. Figure 2.13 shows the new system architecture, where the reference and mobile nodes have identical RF hardware. Figure 2.3 shows the classical approach where the locator hardware is a transmit only device. In order to be portable, and power efficient, the new units are driven with Spartan 6 series FPGAs, and crystal oscillators. The clock prediction scheme first described in [5] allows us to operate without cumbersome atomic clocks. The locator is a two board solution

[6] $20.48 \mu s$ results in a spatial ambiguity of several miles

that combines the functionality of previous locators, as well as on board inertial sensors, and the bi-directional multicarrier software radios. The design of this new PPL revision 3 hardware is thoroughly described in [9].

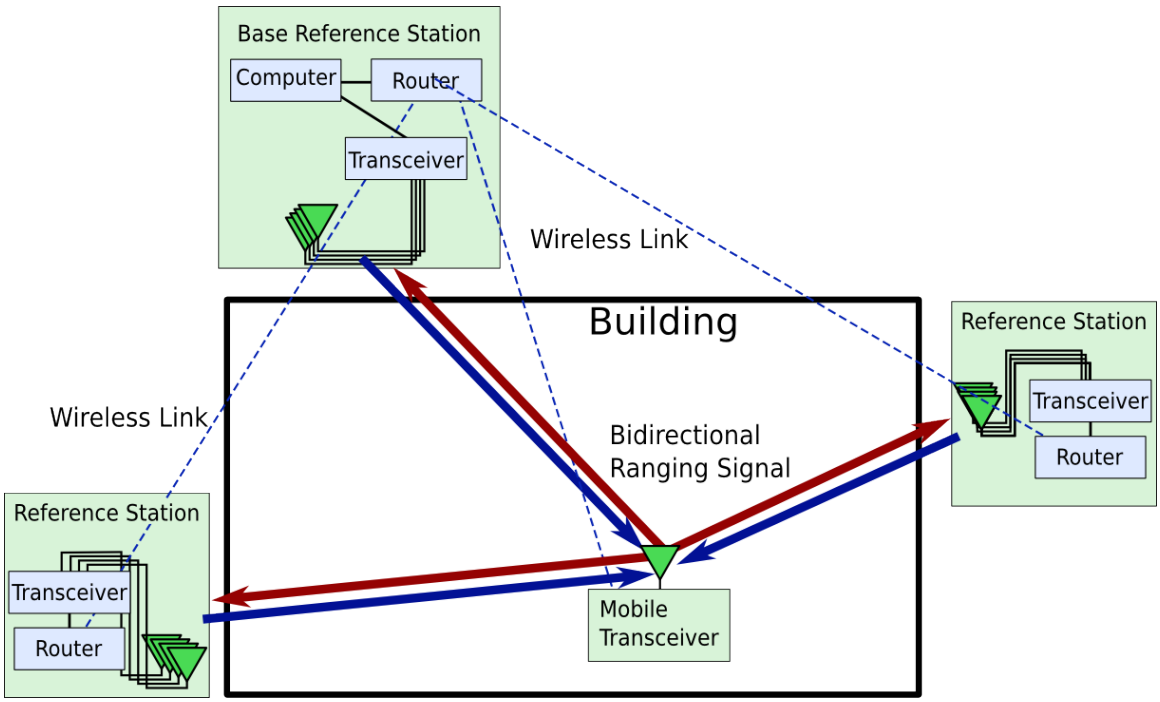

Figure 2.13: TART System Layout [5]

2.5 Typical Field Test

In order to measure the accuracy of our system, we hand-survey our reference antenna locations, as well as mobile truth points on a shared, local, coordinate system. This involves picking a set of points on a floor, and using measuring tapes, and a laser level to obtain the full 3D locations of any additional indoor truth points or outdoor reference points. Once the survey is complete, we set up our reference antennas outside of the house. The antennas are grouped into "arrays" of up to 4 antennas each, corresponding to a given transceiver box. The locator is mounted on a cart, in order to have a consistent known, height for each truth location. Once the transceiver boxes are connected to the base station computer, either over Ethernet or WiFi, we "coarse-sync" all of the boxes. This step ensures that all of the boxes have their internal clocks synchronized sufficiently to have time

offsets less than one signal period. This will be necessary for "fine synchronization" which will complete the synchronization of the clocks.

Once we have coarse synchronization, we are able to start capturing data, which can either be fed into a real time processing algorithm (which performs fine-synchronization in real time), or be logged for post-processing. Typically we will do static captures[7], in which the locator is left at each truth point, and we capture several consecutive symbols of data from the locator. Capturing the data in this manner allows us to get a good idea of what kind of accuracy the system is getting in a given configuration. It also allows us to experiment with different options in post-processing, such as reducing the capture length, or the number of receive antennas to measure their effects on our accuracy. We can also run our real-time capturing system, which processes every symbol captured as it receives it. This allows us to track moving objects, and display their locations in real time, which is ultimately what our system is designed to do. With the current hardware, only σART can be run in this real time mode.

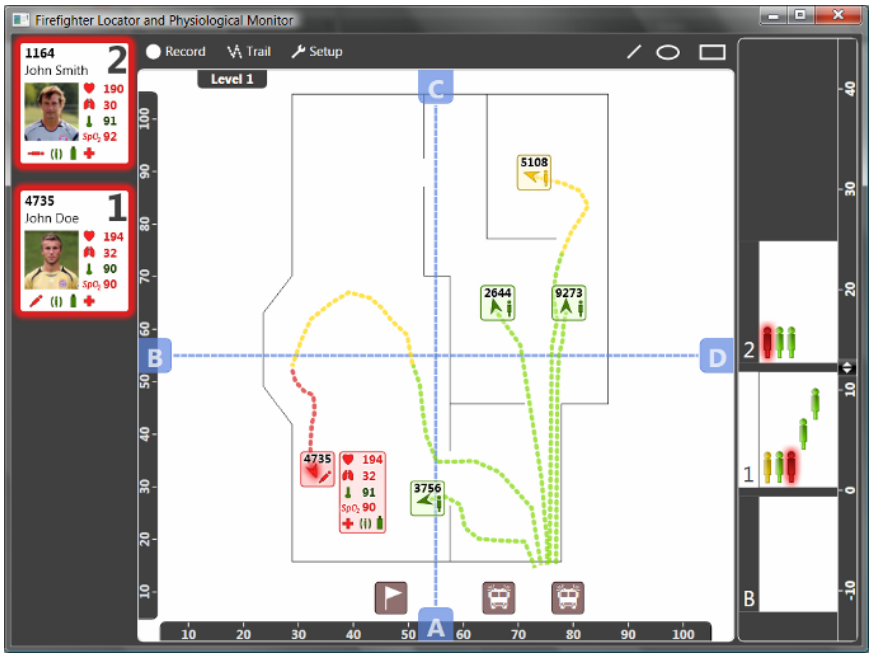

Figure 2.14: PPL Real Time System GUI Designed by QinetiQ North America

Figure 2.14 shows the graphical user interface (GUI) that was created for the PPL system with

[7]Real-time captures allow us to do a more realistic test or demonstration, but are not as useful for testing absolute accuracy with respect to truth locations.

Figure 2.15: Campus Religious Center: Surveying reference antennas in the rear of the building (left), and reference antennas in side and front yards (right)

the help of QinetiQ North America, which shows the locations and tracks of several firefighters. The top down floor plan view is showing the first floor of the structure, on the right hand side of the screen, a representation of the heights of the different responders is shown. The panel on the left hand side shows the ID numbers and photos of the responders, accompanied by real time physiological data taken from the WPI Pulse Oximeter [14] and Foster Miller PSM shirt.

2.5.1 WPI Campus Ministry Center

The WPI Campus Ministry Center is a 3 story wooden house on a typical urban block near the WPI campus. We use this building to evaluate the performance of our system in a typical residential setting. The floor plan seen on the GUI in Figure 2.14 is the first floor of the Campus Religious Center. Several photographs from field tests the Campus Religious Center taken between 2008 and 2010 are shown in Figures 2.15 and 2.16. As seen in the photographs, the building is a wooden house with a wooden frame, with significant impediments to RF in and around the kitchen and the chimney. Structures like this are extremely hazardous to firefighters, and the problem of performing location estimation in such buildings is not trivial. The Campus Religious Center presents our best opportunity to work in a realistic environment for fire fighting incidents. We have also used this building over the past several years to execute realistic firefighting scenarios, and showcase the latest location and tracking technologies being developed by ourselves and others, at our annual workshop.

Figure 2.16: Campus Religious Center: base station set up in living room during a winter field test (right), static capture in the kitchen (left)

2.5.2 WPI Atwater Kent Laboratories

WPI Atwater Kent Laboratories is the Electrical and Computer Engineering building on the WPI campus. The building was originally a brick structure, that was mostly open in the interior, but has undergone several major renovations, including a glass and steel addition on the northern side (see Figure 2.17). Today, the interior of both the addition, and the original structure are filled with offices, classrooms, and labs. The interior walls are drywall over steel studs, as seen in Figure 2.19 and there are many RF obstacles, including: steel beams, electrical conduit, communication wiring, HVAC systems and duct work, fire sprinkler systems, metal light fixtures, chalk boards, fire glass, and metal furniture in every room. All ceilings and floors are also built on metal platforms, as seen in Figure 2.18. This environment is perfect for evaluating our system for performance in medium scale commercial applications.

Figure 2.17: Exterior Photograph of Atwater Kent Laboratories showing the glass and steel addition

Figure 2.18: Photograph showing the metal ceiling and floor decks in the Atwater Kent building

Figure 2.19: Photograph showing the exposed metal studs during a renovation in Atwater Kent Laboratories

2.5.3 Fort Devens Massachusetts

Fort Devens airfield, in Ayer Massachusetts, is a large former airfield, complete with runway. In 2009 we used this site to test the operational range of our system, with our second generation locators, operating in high power mode. This location allows us to test with receiving antennas 1KM away, and was the proving ground for our wireless communications capability. Received signal data, for location and synchronization, was sent over off the shelf wireless routers (IEEE 802.11g), running custom firmware, to a base station at one corner of the airfield. Figure 4.7 shows the Ft. Devens test setup from the perspective of the base station computer. This test site and mission was also used to test and improve our processing speed, by challenging us to track vehicles, moving at speeds up to 40 MPH.

Chapter 3

Theory of Bayesian Information Fusion

This chapter introduces, and provides the mathematical support for, the PPL Bayesian Fusion Algorithm. The general idea of any Bayesian estimation method is to use probability theory to estimate the statistics of a parameter using any available and presumed information about the state of that parameter. This is different from what some mathematicians call the Frequentist's view of statistics, in which making use of less than evidenciary prior information would be improper. The Frequentist's view of statistical inference attempts to estimate the value of an unknown parameter, while a Bayesian approach treats the parameter as a random variable and attempts to estimate its statistical distribution. Frequentists still use models with unknown parameters, and elements of probability theory, including Bayes Theorem. Both views are widely used, and neither one is considered to be strictly better than the other[1]. However, the Frequentist view does not allow for the incorporation of presumed prior knowledge, as it is said to bias the outcomes based on the beliefs of a researcher. The Bayesian approach accounts for the known uncertainty in the parameter by selecting a prior distribution which is sufficiently vague, such that with enough data samples the prior has less impact on the final result. The Bayesian approach promotes the notion that the computed posterior distribution gives us a good picture of the uncertainty of the estimate, rather than just an estimate of a parameter [11, 18, 8]. One of the key dividing points between the two views is whether or not the a posteriori probability density, derived from the data and the a priori density, has a real meaning [2].

[1] Fundamentalist adherents will, of course, claim that their view is superior.

[2] The Bayesian would argue that it does.

In our case we know beforehand where we expect to find firefighters. With our prior knowledge we can greatly reduce the search space for solutions to a discrete number of floor heights constrained by the perimeter of an incident site. This aids accuracy and reduces computation time compared to scanning the entire space of our signal aliasing cell (roughly $8,000,000\,m^3$). Furthermore, our individual sources of information about the location are analogous to likelihood functions, so it makes sense to keep our final solution in terms of a likelihood function, representing our posteriori information, rather than just producing a single answer. By producing the entire likelihood function it allows us to see where the probability of errors is large, and in what direction the errors are most likely to move the solution, as well as clearly showing possible multiple solutions.

Our algorithm is designed to utilize information from both the σART and TART metric functions, as well as additional sensor information, such as barometric height estimation or inertial navigation information, to improve the accuracy of our location system. This approach also improves the robustness of the location solution to an error introduced by any one information source, as long as the errors are not strongly correlated with those of the other sources. [10, 24, 20]. Currently, we have the ability to capture TDOA like data for RF based positioning, TOA like data for RF based ranging, barometric data for height estimation, and six axis inertial measurement unit (IMU) data. The IMU data, while being a useful inclusion for the purposes of enhancing location solutions in high multipath environments, goes beyond the scope of this thesis.

3.1 Bayesian Statistics

Bayesian Statistics is the branch of statistics that attempts to characterize problems where observed data, $\mathbf{D} = [D_1, D_2, \ldots, D_N]$, is used in conjunction with a prior probability distribution, based on information, I, that is known a priori. Equation (3.1) is the familiar statement of Bayes' Theorem for events A and B.

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \tag{3.1}$$

In our case, we are applying Bayesian statistics to estimate three parameters, the x, y and z coordinates of our locator, by finding the location where the posterior likelihood function is maximized. The prior probability distribution for our location solution, where we assume that the locator lies

somewhere on a bounded 'scan grid', is given by (3.2).

$$P(\mathbf{x}|I) = \begin{cases} k & , \mathbf{x} \in \text{scangrid} \\ 0 & , \text{otherwise} \end{cases} \tag{3.2}$$

Where $\mathbf{x}$ is used to represent the set of parameters:$\{x, y, z\}$, and k is a constant, chosen to ensure the distribution integrates to unity. When we consider the measured data, we obtain the posterior PDF, given by (3.3).

$$P(\mathbf{x}|\mathbf{D}, I) = \frac{P(\mathbf{D}|\mathbf{x}, I) \times P(\mathbf{x}|I)}{P(\mathbf{D}|I)}$$

$$P(\mathbf{x}|\mathbf{D}, I) \propto P(\mathbf{D}|\mathbf{x}, I) \times P(\mathbf{x}|I) \tag{3.3}$$

The metric function for a datum corresponds to the term $P(\mathbf{D}|\mathbf{x}, I)$ of (3.3), which describes the probability that the metric function is maximized at each point in the scan grid, given the correct location. This term is multiplied by (3.2)

Our three data[3] are assumed to have uncorrelated error characteristics, as shown in Section 3.2. Under a further assumption of Gaussianity, they may be taken as independent. Since they will be treated as independent, we can apply Bayes' Theorem, as well as the the identity: $P(A|B, C) = P(A|C)$ when A and B are independent.

$$P(D_\sigma, D_T, D_b|\mathbf{x}, I) \propto P(D_\sigma, D_T|D_b, \mathbf{x}, I) \times P(D_b|\mathbf{x}, I) \tag{3.4}$$

$$P(D_\sigma, D_T|D_b, \mathbf{x}, I) \propto P(D_\sigma, D_T|\mathbf{x}, I) \tag{3.5}$$

$$P(D_\sigma, D_T|\mathbf{x}, I) \propto P(D_\sigma|D_T, \mathbf{x}, I) \times P(D_T|\mathbf{x}, I) \tag{3.6}$$

$$P(D_\sigma|D_T, \mathbf{x}, I) \propto P(D_\sigma|\mathbf{x}, I) \tag{3.7}$$

$$\therefore P(D_\sigma, D_T, D_b|\mathbf{x}, I) \propto P(D_\sigma|\mathbf{x}, I) \times P(D_T|\mathbf{x}, I) \times P(D_b|\mathbf{x}, I) \tag{3.8}$$

This result will be applied in Section 3.3.

3.1.1 The Debate Between Bayesians and Frequentists

To see the fundamental difference between the Bayesian and the Frequentist, consider the following non-probabilistic example, first proposed by Professor Cyganski as an illustration of the

[3]The σART metric, TART metric, and barometric height estimation.

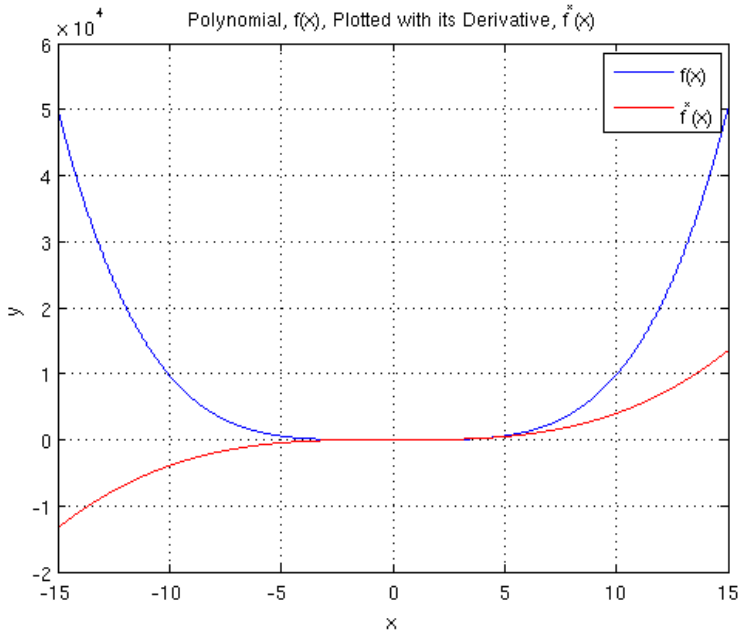

Figure 3.1: Plot of a polynomial function (blue) with its derivative (red)

core argument taken to an extreme. Let's say that fictitious mathematicians have applied the notions of Bayesians and Frequentists to polynomials. Let's call them P-Bayesians and P-Frequentists. A P-Bayesian and a P-Frequentist are both presented with the same problem: given a polynomial function of x over an interval where it has a zero crossing, like the function shown in Figure 3.1, *estimate* the location of the zero crossing. Any introductory calculus text will tell you to apply Newton's method to find an appropriate solution to this problem. Given $f(x)$ and its derivative, $f'(x)$, the following recursion relation will converge to the correct value:

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}$$

We will consider each evaluation to be a new piece of experimental data. This method is extremely effective, and common knowledge to anyone with sufficient background in mathematics, but there is one small problem. In order for this recursive algorithm to do anything, we need an initial value, x_0. The P-Bayesian will choose what they believe to be a valid value for x_0, based on their experiences with similar problems, and perhaps some intuition. This is where the P-Frequentist objects, choosing an initial point will bias the outcome based on the beliefs and intuition of the P-Bayesian! After only one of two peices of experimental data (the outcomes of one or two iterations of New-

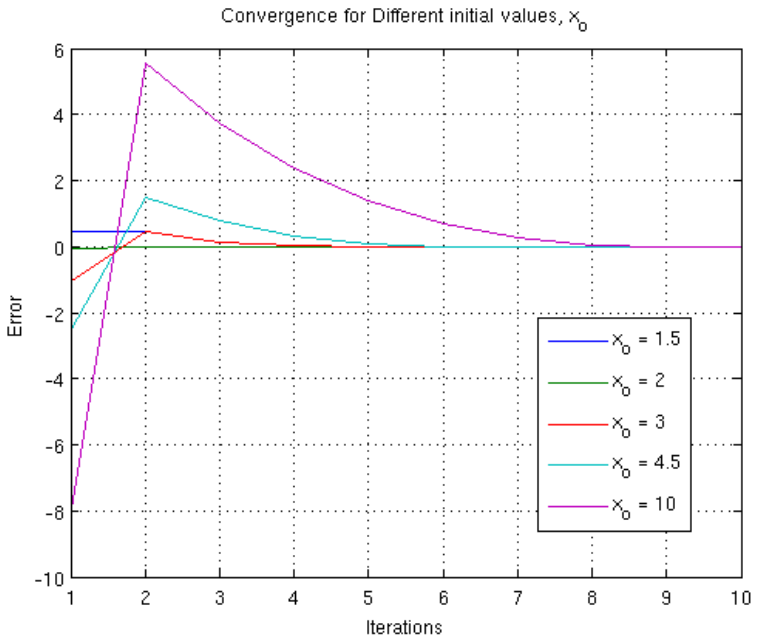

Figure 3.2: Convergence of Newton's Method to analytic result for different initial points

ton's Method). This is not sound science, and cannot possibly be a valid approach. While the P-Frequentist ponders an alternative solution, the P-Bayesian begins computing their solution based on the chosen initial value. The results of 10 iterations with 5 different choices for X_0 are shown in Figure 3.2. The error in this plot is measured as the difference between the approximation and the exact solution (1.95458 according to Wolfram Alpha). A good initial guess can certainly make the approximation converge faster, as is the case for the red, blue, and green curves. When the initial guess was not as good, the recursive method still pulled the solution to the correct value within 10 iterations, at which point all of the solutions converge within the precision of the given final answer.

In the parallel of a Bayesian estimation problem, the initial guess would be equivalent to the a priori distribution assigned by the analyst, and the data would shape the posterior distribution in much the same way that Newton's Method applied repeatedly successively, and successfully, approximated the zero crossing for the function of interest.

Let's assume that the P-Frequentist has now thrown up his or her hands and decided to choose a starting point, as long as it is done in a scientifically valid way (we will not give a construction for such a method here). What answer will the P-Frequentist give to the problem? After how many iterations will the solution be deemed acceptable? While there may be criteria for a solution's correct-

ness based on the problem statement, the P-Frequentest will simply iterate until the error is within the acceptable tolerance. The P-Bayesian, on the other hand, will look at the intermediate results, and notice that sometimes the solution converges in much fewer than 10 iterations. Furthermore, they may use knowledge gleaned from this experiment to better choose initial values in the future. An inference that a P-Bayesian might draw from this experiment would be that if computation is expensive then it is a good idea to pick a point that is believed to be close to the correct answer. If, on the other hand, computation is cheap, then any guess that is reasonably close to the correct answer will converge in finite time to an arbitrary closeness, with the initial guess having very little to do with the final outcome (other than the fact that it enabled the analysis to be performed in the first place). For a great example of Bayesian estimation with varied priors, see [24].

To summarize, both Bayesians and Frequentists are interested in using data to estimate unknown parameters. The Frequentist considers the parameters to unknown quantities whose values must be estimated, while the Bayesian is interested in the probabilistic distribution of the parameters. Both types of statistician use probability theory, including Bayes' Theorem, but the Frequentist will only use probabilities that are based on the concept of relative frequency on repeated trials in a probability space, while the Bayesian is allowed to assign distributions a priori based on assumed knowledge of the specific problem[11, 24].

3.2 Weak Correlation of σART and TART Errors

In order to justify the assumptions made in the implementation of our fusion algorithm, we need a basis for saying that the σART metric function and TART metric function are representative of two independent random variables. Since both algorithms yield estimates of the same parameter (position), using the same data, the σART and TART metric functions for a given received signal data matrix should both be maximized at the correct location in the scan grid. The two algorithms are internally similar up to the re-phasing step, covered in Section 2.2.2. After this point the σART algorithm evaluates an SVD-based metric, which ignores the unknown time delays seen by all antennas. The TART algorithm solves for the unknown absolute time delays between the locator and all reference antennas, followed by evaluating a much simpler summation based metric.

The likelihood that a given point in a metric function is it's maximum value, given a true location

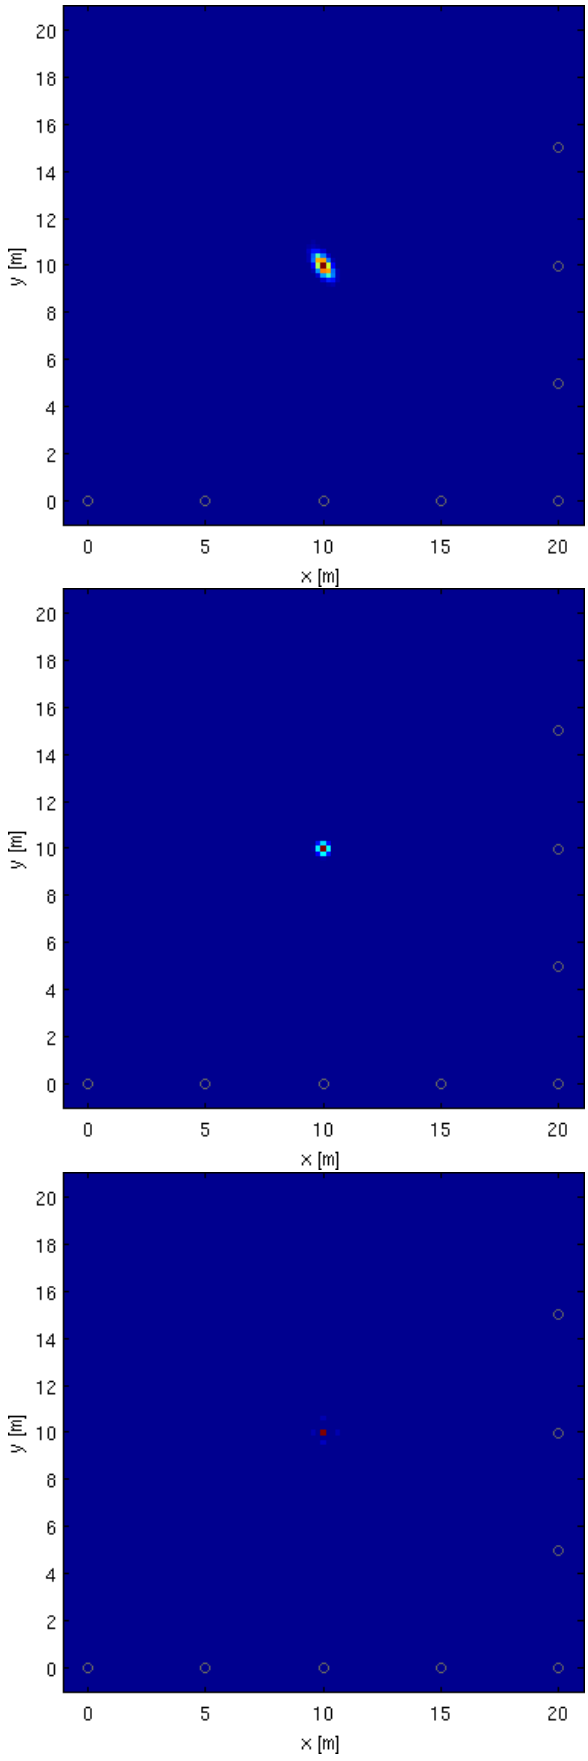

Figure 3.3: σART (top), TART (middle) and Bayesian Fusion Algorith (bottom) results for 10,000 Monte Carlo trials -6dB SNR, and 8 receiving antennas on two sides

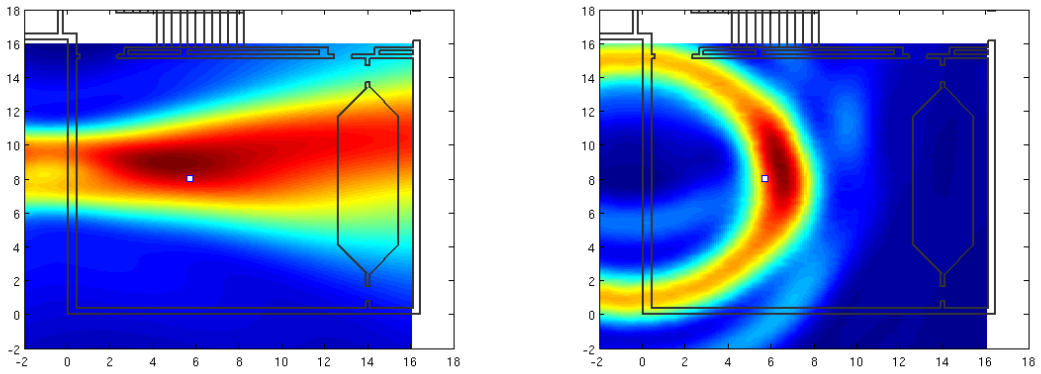

(a) σART (b) TART

Figure 3.4: σART vs. TART - AK Ladders

is proportional to the metric function itself in the case of σART and TART. When the RF channel is very clean, the two metric functions look extremely similar, both are peaked around the same point. The estimates in this case have similar distributions, having nearly identical means, and small variance. This is illustrated in Figure 3.2, where repeated trials produced an error distribution with the expected skew in σART and TART results, with the variance of the fused result being much smaller than either of the components. In situations with reduced SNR, when suffering from geometric dilution of precision (GDOP), or high multipath, the two resulting metric functions typically look very different. If we think of the normalized likelihood functions described in Sections 2.2.2 and 2.3.2, these differences can be explained as the errors from the two algorithms having orthogonally skewed probability distributions. This basic assumption has been tested in both simulation and real field tests, and holds true in most cases.

To verify that the errors of σART and TART are not highly correlated, 50,000 Monte Carlo trials were conducted. These trials were the same to those of Figure , but the SNR was varied from 0dB to -12dB, with 10,000 trials conducted at each noise level. From these errors, we calculated correlation coefficients, $\sigma_{x,y}$, which range from 0 (totally uncorrelated) to 1 (completely correlated). These coefficients are given in a matrix with the form of Equation (3.9).

$$\begin{pmatrix} \sigma_{x,x} & \sigma_{x,y} \\ \sigma_{y,x} & \sigma_{y,y} \end{pmatrix} \tag{3.9}$$

Where $\sigma_{x,y}$ is the correlation coefficient between the σART errors in the x direction and the TART

errors in the y direction (the first subscript corresponds to σART, and the second subscript corresponds to TART). The results of 10,000 Monte Carlo trials at each SNR are listed below:

$$0\text{dB} = \begin{pmatrix} 0.64 & 0.00 \\ 0.00 & 0.49 \end{pmatrix}$$

$$-3\text{dB} = \begin{pmatrix} 0.60 & 0.00 \\ 0.00 & 0.45 \end{pmatrix}$$

$$-6\text{dB} = \begin{pmatrix} 0.53 & 0.00 \\ 0.01 & 0.42 \end{pmatrix}$$

$$-9\text{dB} = \begin{pmatrix} 0.40 & 0.01 \\ 0.01 & 0.11 \end{pmatrix}$$

$$-12\text{dB} = \begin{pmatrix} 0.03 & 0.00 \\ 0.00 & 0.02 \end{pmatrix}$$

As the SNR decreases, the errors grow, and the correlation coefficients decrease. Figure 3.2 shows that the errors are still relatively small at -6dB SNR, yet the correlation between σART and TART errors for a given axis is close to 50% and decreases with SNR. These trials, as well as data taken from field tests supports the assumption that the errors for the two algorithms are not highly correlated.

Figure 3.4 shows the typical shape of the σART and TART metric functions from a field test in the Atwater Kent Laboratories. This was a situation with a very severe GDOP, and a difficult multipath channel. The image on the left shows the metric function of the σART algorithm. The white dots represent the receiving antennas, mounted on four ladders against the outside of the building. The white square is a transmitter location inside of a classroom. The elongated hexagon towards the right of the image is a lowered portion of the room with a podium and several double-hung blackboards. The colors of the metric function represent its value, with blues corresponding to low values, and reds corresponding to high values. The σART metric function is 'blurred' along the X axis, in the direction away from the plane of antennas. Meanwhile, the image on the right shows the TART metric evaluated on the same set of data. Instead of 'blurring' along a hyperbola,

as we saw in the σART metric, the TART metric is ambiguous along the radius of a circle, with the reference antennas that are giving the strongest contribution at its center. This type of circular symmetry is what one would expect from a TOA like positioning algorithm. While neither metric is maximized at the correct location, they are both very near maximum in the neighborhood of the correct point. Furthermore, in this particular case it is clear that the TART algorithm has a much better estimation of the X coordinate of the locator, while the σART algorithm has a better estimate of the Y coordinate. These orthogonal error characteristics and the demonstrated strength of the assumption of independence between σART and TART promote the notion that the Bayesian Fusion methods described above will yield performance benefits.

It is obviously true that the barometric height estimates also have errors that one would expect to be uncorrelated with the σART and TART errors. This sensor measures pressure, which is an entirely different physical parameter than RF signal characteristics. For this reason the barometric sensor is extremely useful for this type of information fusion, as it is truly an independent measure.

3.3 PPL Bayesian Fusion Algorithm

In order to build an estimator for the position of our locator, it is first necessary to construct probability density functions (PDF), or likelihood functions, that correspond to the received data, as well as the prior information. In our case we will start by assuming a uniform prior, for our estimated position, which is uniform on the area contained in our search space[4]. This relatively simple assumption can be enhanced if any other information about the structure is known, for example the inter-floor spacing, or the location of stairs/elevators. Once all of the prior information is taken into account, we then incorporate our data into the solution via the Bayesian estimation process described in Section 3.1.

We will assume here, that we are given the PDFs generated from the barometric, σART , and TART data. The process of generating these PDFs from raw data will be covered in the following chapter. To estimate the parameters, x, y, z, which are the X, Y, and Z coordinates of our locator, we discretize the posterior distribution into a number of points, which lie on planes that are stacked

[4]Although we are doing 3D location, the word 'area' is used here because we are scanning 2D slices of a 3D space.

along the Z-Axis. For every point in the discrete scan grid, we evaluate:

$$P(D_\sigma, D_T, D_b|\mathbf{x}, I) \times P(\mathbf{X}|I) \tag{3.10}$$

for

$$\mathbf{x} = x, y, z$$

The point, x, y, z, where this metric is maximized corresponds to the point in space where the locator is estimated to be, based on the available data.

The Bayesian Fusion algorithm was developed in order to maintain the accuracy of the PPL system in situations where we have poor antenna coverage, or very low SNR. Situations in which only one side of a building is covered are of particular interest because of the orthogonally skewed distributions of the σART and TART likelihood functions. While all three sources of data affect the final 3D solution, it is clear that in the case of a planar array of antennas on the y, z plane (where the z axis corresponds to height), σART provides the best data about the x coordinate, TART is most accurate in the y coordinate, and the barometric information is the best source for z information. The Bayesian Fusion Algorithm is not altered for this situation, but the components of the resulting distribution are largely shaped by the aforementioned data sources in their respective directions. While the barometric sensor never gives us data regarding the x, y position of the locator, knowing what 2D plane to scan in can greatly reduce the ambiguity in the σART and TART data sources. Being able to work in less ideal settings is a critical part of having the capability to deploy the system in a realistic setting, where time and space for system deployment are tightly constrained.

Chapter 4

Rapid Deployment Capability

Ideally, the PPL system would not require the addition of extra personnel, or setup time, to a fire fighting effort. Until now, we have relied on: hand-surveyed antenna locations, extensive cabling between remote locations and the base station, as well as antennas being deployed on all sides of a building. This effort takes a team of several people close to an hour to complete. This chapter will detail the newly acquired rapid deployment capabilities for the PPL system, as well as the trade-offs associated with these new methods.

In the past year, we have made several advances towards having a system that is automatically deployable. Mounting antennas on ladders, at known distances from each other allows us to quickly deploy large numbers of antennas, with relatively little information needed to obtain their locations in three dimensional space. Wireless communication, and new synchronization hardware and software techniques, have allowed us to do everything that our previous system does, without additional wires needed for data transmission, or clock distribution. Most of the setup time is spent surveying the locations of the receiving antennas; this entire step is supplemented by the hardware with techniques for Geometric Autoconfiguration (GAC).

4.1 Fixed Antenna Arrays

Mounting our patch-style ranging antennas on aluminum extension ladders, as seen in Figure 4.1, helps us achieve our goal of having a rapidly deployable system. There are several advantages to our ladder mounted antennas,. The first is that the setup time, on site, is reduced significantly

by the fact that the antennas are attached to the ladders ahead of time. The second advantage to this configuration is that the inter-antenna distances for antennas on a single ladder are known. This greatly simplifies any site-surveying methods, and greatly benefits the accuracy of GAC techniques through the inclusion of fixed constraints. The ladders also enable four antennas to be deployed in the same amount of time that it would take to deploy a single antenna in the classical configuration.

Since the ladders that we use are made of aluminum, we currently use folded-patch antennas, with the ground plane mounted against the ladder, on the side facing into the building. If we were to use antennas without ground planes, the ladder would effect the performance of the antenna, in unpredictable, and potentially detrimental ways. Using the patch antennas also rejects any noise or reflections that are generated from behind the antennas.

The ladder approach is a first step towards a rapidly-deployable antenna mount. Several logistical problems still need to be addressed. First, the cables still need to be plugged in once the ladder is put up. Also, the presence of the patches makes climbing the ladder much more difficult, and the ladders should really have safety ropes to keep them from falling over. All of this means that additional setup time is required. We currently have, under development, a new antenna concept that will eliminate many of these problems, but ultimately, the antennas will need to be integrated with the fire truck, police cars, or other vehicles. This will completely remove the need for plugging in cables and setting up antennas. In theory, a button will simply be pressed, and the network of transceivers will perform GAC to calibrate themselves and then begin the full personnel tracking function.

4.1.1 Accuracy Considerations

Preliminary RF tests have shown that a geometric dilution of precision (GDOP) is introduced when the ladder mounted antennas are all deployed along one side of a building. This is because the antennas are arranged in a nearly co-planar constellation. The effects of this GDOP are worst when the σART solution is computed, as the σART method has very poor resolution in the direction perpendicular to the plane. This is because the TDOA-like metric is related to the hyperbolic curves of traditional TDOA estimators. The TDOA between the elements on the plane changes much more slowly than the distance between the transmitter and the array. As the distance grows larger, the GDOP becomes worse as the plane becomes more point-like. The TART solution is also adversely

affected by this arrangement. The TART solution shares many features with a TOA solution in which the point where all of the spheres intersect is the correct point. If two antennas are located at the same point, their spheres will intersect at all of the points on their surfaces. If the distance between the reference antennas is large compared with the distance, these spheres will intersect at a single point. In our case, the distance is not large enough to produce a well defined intersection.

4.2 Utilization of Additional Sensors

In order to maintain a high degree of accuracy even in situations with high geometric dilutions of precision, we needed to incorporate more information into our position estimation algorithms. This information includes 1D ranging, barometric height estimation, and inertial navigation supplementation [4]. The 1D ranging and barometric height estimation are detailed below. While the inertial navigation can also be fused with our RF and barometric data, it is beyond the scope of this thesis. For a proper treatment of this topic with regards to the PPL project see[22].

4.2.1 1D Ranging

Recently, the work by V. Amendolare [5] has resulted in system modifications that allow us to capture TOA-like data between pairs of transceiver boxes. This makes it possible to get individual 1D ranging estimates, which are useful in situations where antenna geometry constraints produce a geometric dilution of precision in one or more dimensions [10]. The σART metric, detailed in Section 2.2.2, ignores constant time offsets, as it is a fundamentally TDOA-like approach. While this proves to be advantageous for synchronization [1], the absolute distance information is being thrown away. Using the transceiver box as a 'locator' we are able to use the TART algorithm, described in Section 2.3 to extract the 1D range between the 'locator' and the transceiver units, in addition to evaluating the traditional σART metric. This 1D ranging information is crucial in our rapid deployment setup, where we may only have antennas on one side of a building (see Figure 4.1). These antennas tend to be nearly co-planar, which results in a high GDOP in the direction pointing out of the plane. In the case presented in [10] the building in question was a wooden house, which meant that relatively accurate RF 1D ranging should be possible. The result of fusing the 1D information with the σART information yielded a great improvement of accuracy

Figure 4.1: Rapid Deployment Setup

along the X direction. The error vector plots from this test are shown in Section 6.1.

4.2.2 Barometric Height Estimation

Our barometric sensor, the Freescale Semiconductor MP3H6115A, enables us to collect data on the height of the locator, that is completely independent of any of our RF-based information. Furthermore, the height estimates from the barometric sensor can be fed forward into our location software, and reduce our exhaustive 3D σART or TART scan to a single 2D slice. This speeds up computation, as well as removing any Z ambiguity in the solution space. Height information has been cited as the most important thing that firefighters need to know when searching for lost personnel. Knowing the correct floor of a building is crucial; this means that height errors should never exceed $\pm\frac{1}{2}$floor[1].

In order to accurately estimate height, two effects need to be ameliorated. The first is the time-varying nature of the natural atmospheric pressure. This is easily compensated for by having a reference unit, at a known height, and calculating the pressure difference between the two units.

[1] $\approx$1.5 meters

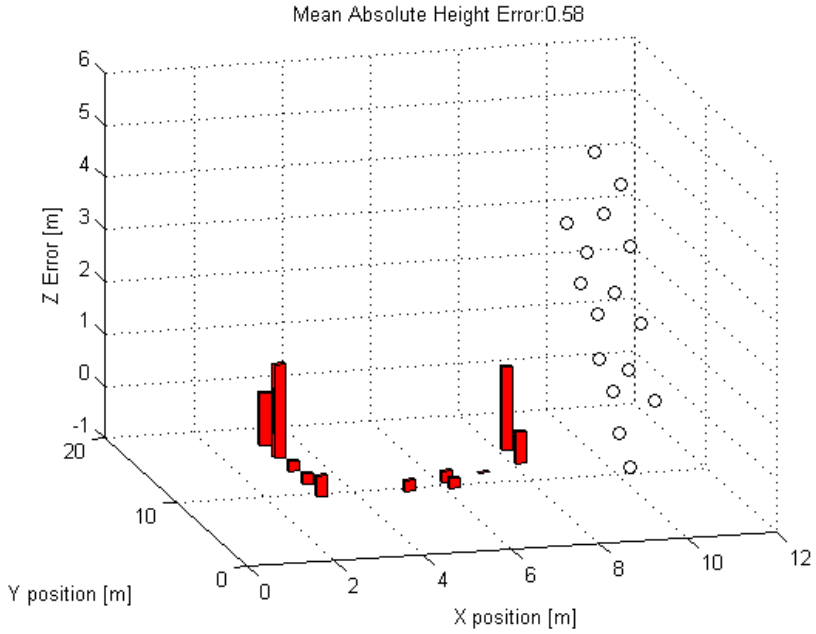

Figure 4.2: Barometric Height Estimate Errors From Campus Religious Center

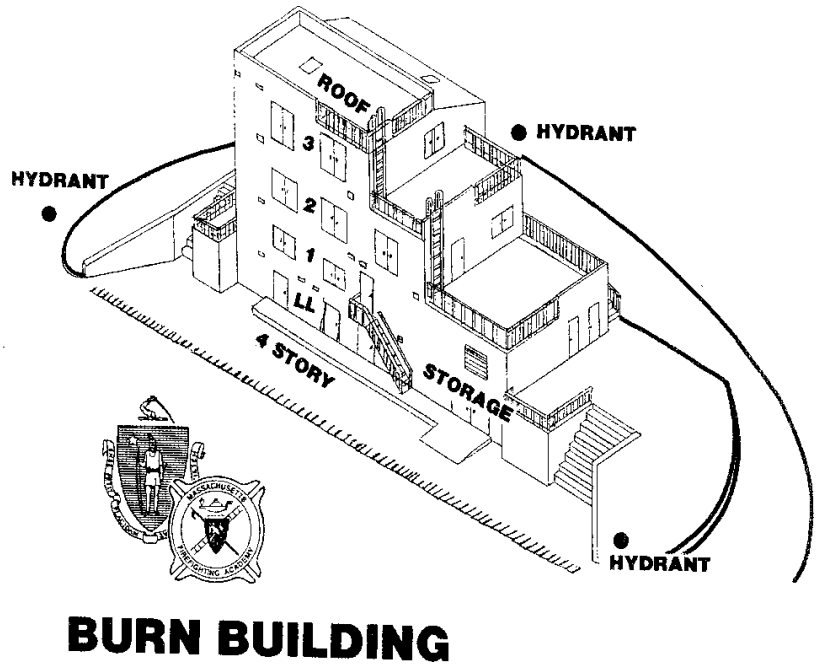

Figure 4.3: Sketch of Stow Burn Building, Courtesy of Massachusetts Firefighting Academy

Figure 4.4: Stow Test: MFA Stow Burn Building, Outdoor Reference/Base Station, Indoor Sensor, Flames from Fire Room

The relationship between pressure and elevation is primarily dependent on absolute altitude, which would be an issue if the system were calibrated at sea level, and deployed at a much higher altitude. The change in air pressure due to height is approximately $-12\frac{Pa}{m}$ at sea level. The second effect that we must mitigate is the dependence of the output of our barometer on the temperature of the unit. We have used the reading of the data-channel board temperature sensor to calibrate the units, under the assumption that the temperature reading was indicative of the barometric sensor's temperature. The temperature effects are modeled by a constant bias, as well as by a scaling factor. To calibrate the units, we took extended data captures on three floors of a building, while capturing temperature data. We also conducted a test in which we varied the temperature of the units, with a space-heater, keeping the units at a constant height. The temperature characteristics of the units were analyzed and we did a linear fit of pressure vs. temperature over our temperature range. To compensate for actual changes in atmospheric pressure, caused by weather, we had a unit that was at the same height, but not near the heat source in order to get only the change in pressure due to temperature.

This calibration effort allowed us to determine the temperature coefficients of each barometric

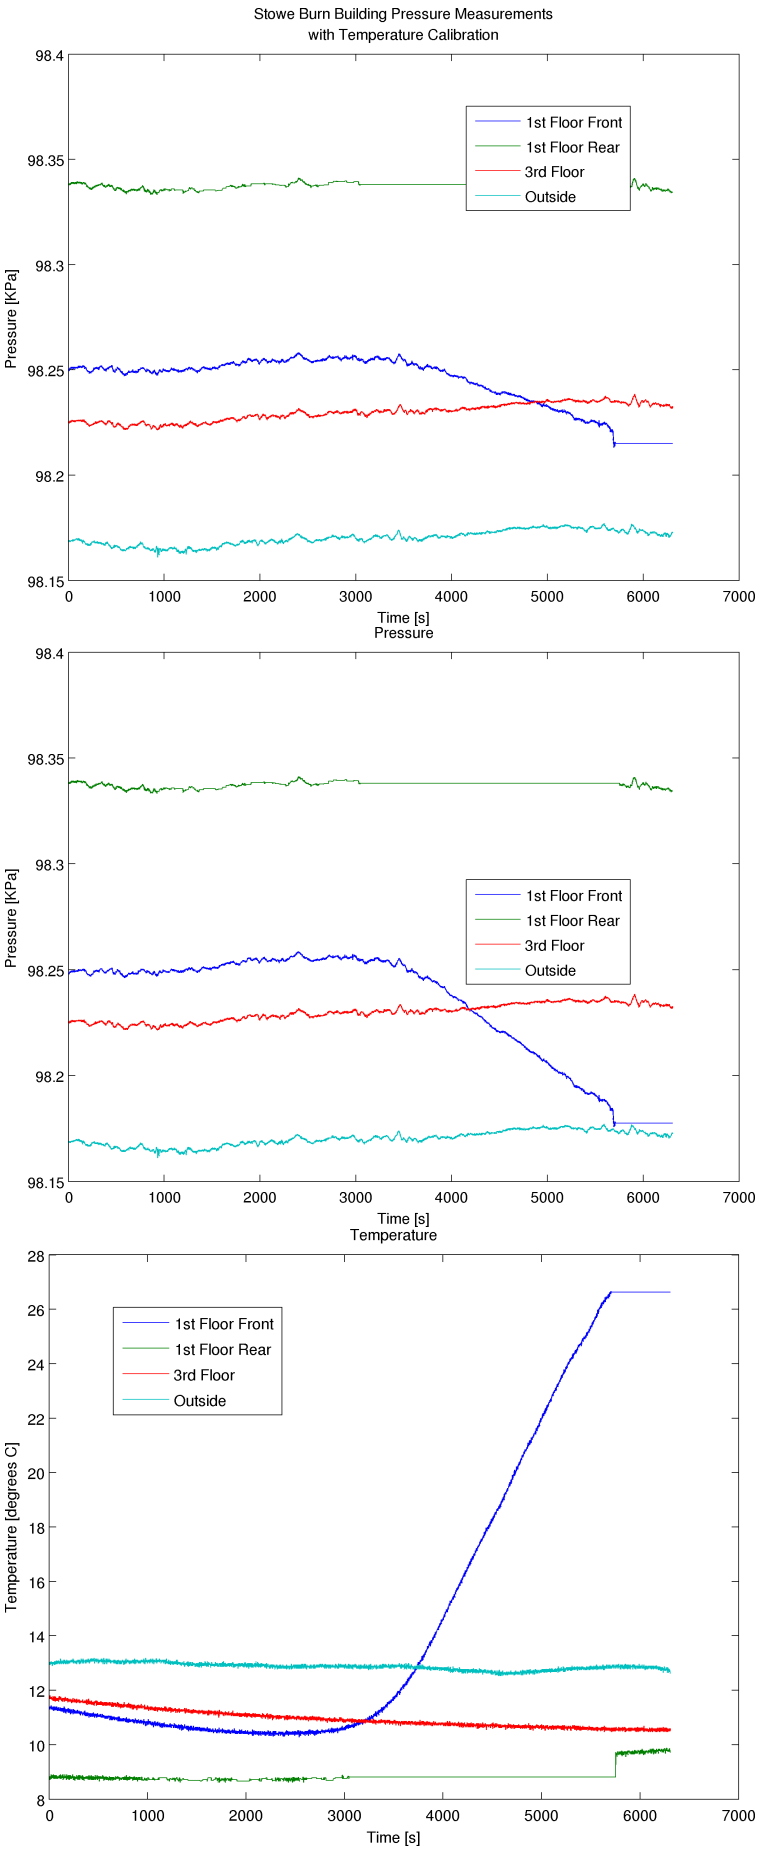

Figure 4.5: Stow Pressure Data: Temperature Compensated Pressure Curves, Uncompensated Pressure Curves, Temperature Curves

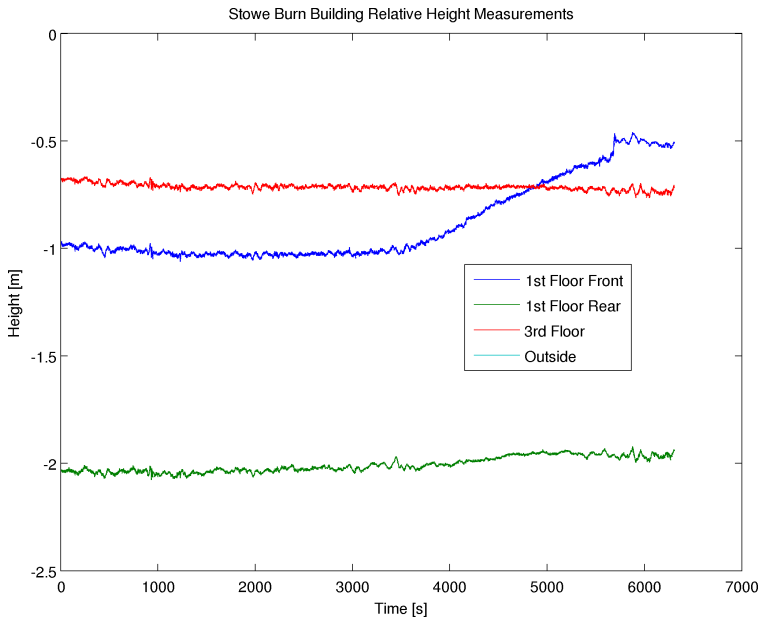

Figure 4.6: Relative Height Measurements from Calibrated Pressure Data Taken at the MFA Stow Burn Building

sensor's pressure offset and scale parameters. From these, individual correction algorithms were developed and applied that allowed us to effectively estimate the height of our locator on all three floors of a wooden house. The height estimates were found to be very robust to temperature changes, as well as events such as doors and windows opening. A plot of the barometric sensor's Z error from the first floor of the WPI Campus Ministry building is shown in Figure 4.2. In this figure, the base of the bars is the location where the unit was placed in the building. The height of the bar is the error in height from the barometric height estimate. The units were later tested at the Stow Burn Building, while a burn test was being performed on the first floor towards the front of the building. Figure 4.3 is a sketch of the burn building at Stow. We placed 3 pressure units inside the building, and 1 unit on the outside for a pressure reference. The unit labeled "First Floor Front" in Figure 4.5 was located closest to the fire, which is clear in the temperature plot[2]. Although the temperature compensation seems to break down at very high temperatures, the biggest challenge that we faced in this test was the RF environment in and around the concrete and steel burn building. Our pressure data was being relayed to us via our 915MHz. data-link, which had long periods of outages. When we experienced an outage, we chose the last known value of pressure. The effects of this approximation are apparent

[2]The temperature readings are taken from from the Analog Devices ADT7301 temperature sensor, and are shown in degrees Celsius.

as the horizontal segments of the graphs in Figure 4.5. This can be seen very clearly with the green curve, which corresponds to the unit at the rear of the first floor.

The Bayesian Fusion Algorithm has proven to be effective for improving our performance with severe GDOP in previous works [10], where it was also found that the computation time of the algorithm could be decreased with a slight modification. Analysis of the barometric data showed the errors to always be less than ± 1 floor. By updating the prior to only contain a plane for one floor, the σART and TART metrics need only be calculated on a single floor. The relatively small standard deviation in the barometric readings, as well as the assumption of a Gaussian PDF, lead to the barometric height estimate determining the floor in the fused algorithm even when we did not run the more computationally efficient algorithm. When the barometric sensor gives a severe error, the standard deviation of the measurement is still small, this is because errors in barometric readings have a large deterministic component, produced by uncompensated swings in temperature or pressure. It is hoped that improved sensors will mitigate this potential source of error. Fortunately, the X,Y estimates are not greatly perturbed by large Z-errors, in cases where there is poor Z resolution in the RF system. If the antenna geometry permitted more precisely defined RF height estimates, it would make sense instead to do a full 3D scan, and give the barometric height estimates a Gaussian PDF so that it could be incorporated into the Bayesian estimation process on equal footing with the RF data.

4.3 Wireless Communications

The ability to communicate wirelessly is essential for control of the locator units. The way we address the problem of serving multiple personnel with locators is by using a TDM system in which the locators are told, by the transceivers, when to transmit, and when to stop transmitting. This allows every locator to take advantage of our full bandwidth, as well as duty cycling power intensive RF hardware to conserve battery power. This control protocol uses the 915MHz data channels in the locators as well as the transceivers. Additionally, off the shelf 802.11g routers are employed to facilitate communication between remote transceiver boxes and the base station. Figure 4.7 shows two tripods outfitted with 18 dBi antennas to relay data over long distances, using 802.11g routers. This is especially important when performing location over large areas, or when

transceivers are located on different vehicles, that do not necessarily arrive at a site together. The wireless configuration is shown in Figure 2.3.

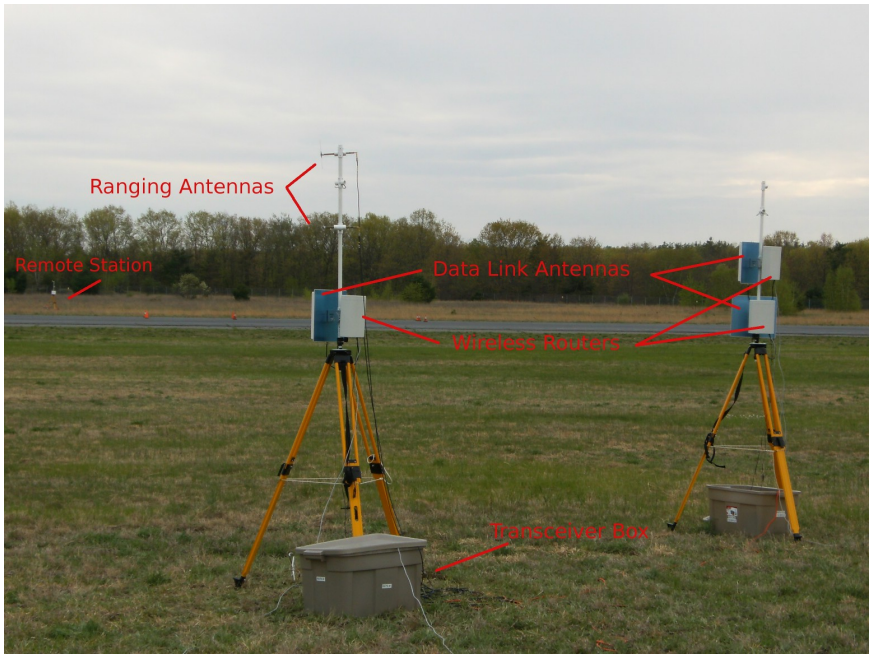

Figure 4.7: Photo of transceiver station and data-link relay

4.4 Geometric Autoconfiguration

The problem of geometric autoconfiguration (GAC) has been extensively studied in a recent thesis [25]. Preliminary results suggest that using GAC would be 3-4 times as inaccurate as using hand surveyed points. While this is expected to improve, there are certainly situations in which this amount of error would be acceptable if it meant being able to set the system up in far less time.

Ultimately, our goal is for GAC to completely remove the need to do hand surveys of antenna locations. The new antennas that we are developing should improve the accuracy of GAC in a ladder-like test setup. Any uncompensated variations in antenna pattern will degrade the performance of GAC; the most extreme case being two antennas that cannot communicate. Having omnidirectional antennas with LOS or NLOS will be the best case scenario for GAC. The advantages that these new antennas have over the patches mounted on ladders, is that they too will be at known distances from one another on a line, but unlike the patches, these antennas would be omnidirectional in the

azimuthal plane, giving the GAC transactions much better inter-array propagation information.

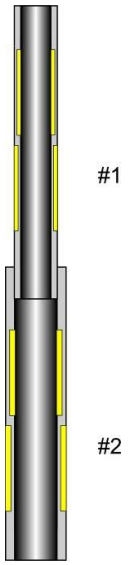

Figure 4.8: Diagram of Telescopic Antenna [16]

Such an antenna design has been completed [16], and is in the prototype stage. A diagram of the antenna configuration is shown in Figure 4.8. Initial tests of the PPL system with this new antenna will be conducted in the coming year. These antennas should outperform even our vertical dipoles, since there is no hardware, such as baluns and cables between the radiator and the far-field. Furthermore, the antennas that are on the same "array" have almost no mutual coupling[3], due to the geometry of the radiation pattern, as well as the inter-antenna spacing. The radiation pattern plots for the E Plane (along the mast), and H-Plane (centered on mast) are shown in Figure 4.9. While being omni-directional in the H-Plane, the null in the center of the E-Plane is responsible for the low inter element coupling for vertically stacked elements.

[3]The coupling predicted by [16] is -30dB.

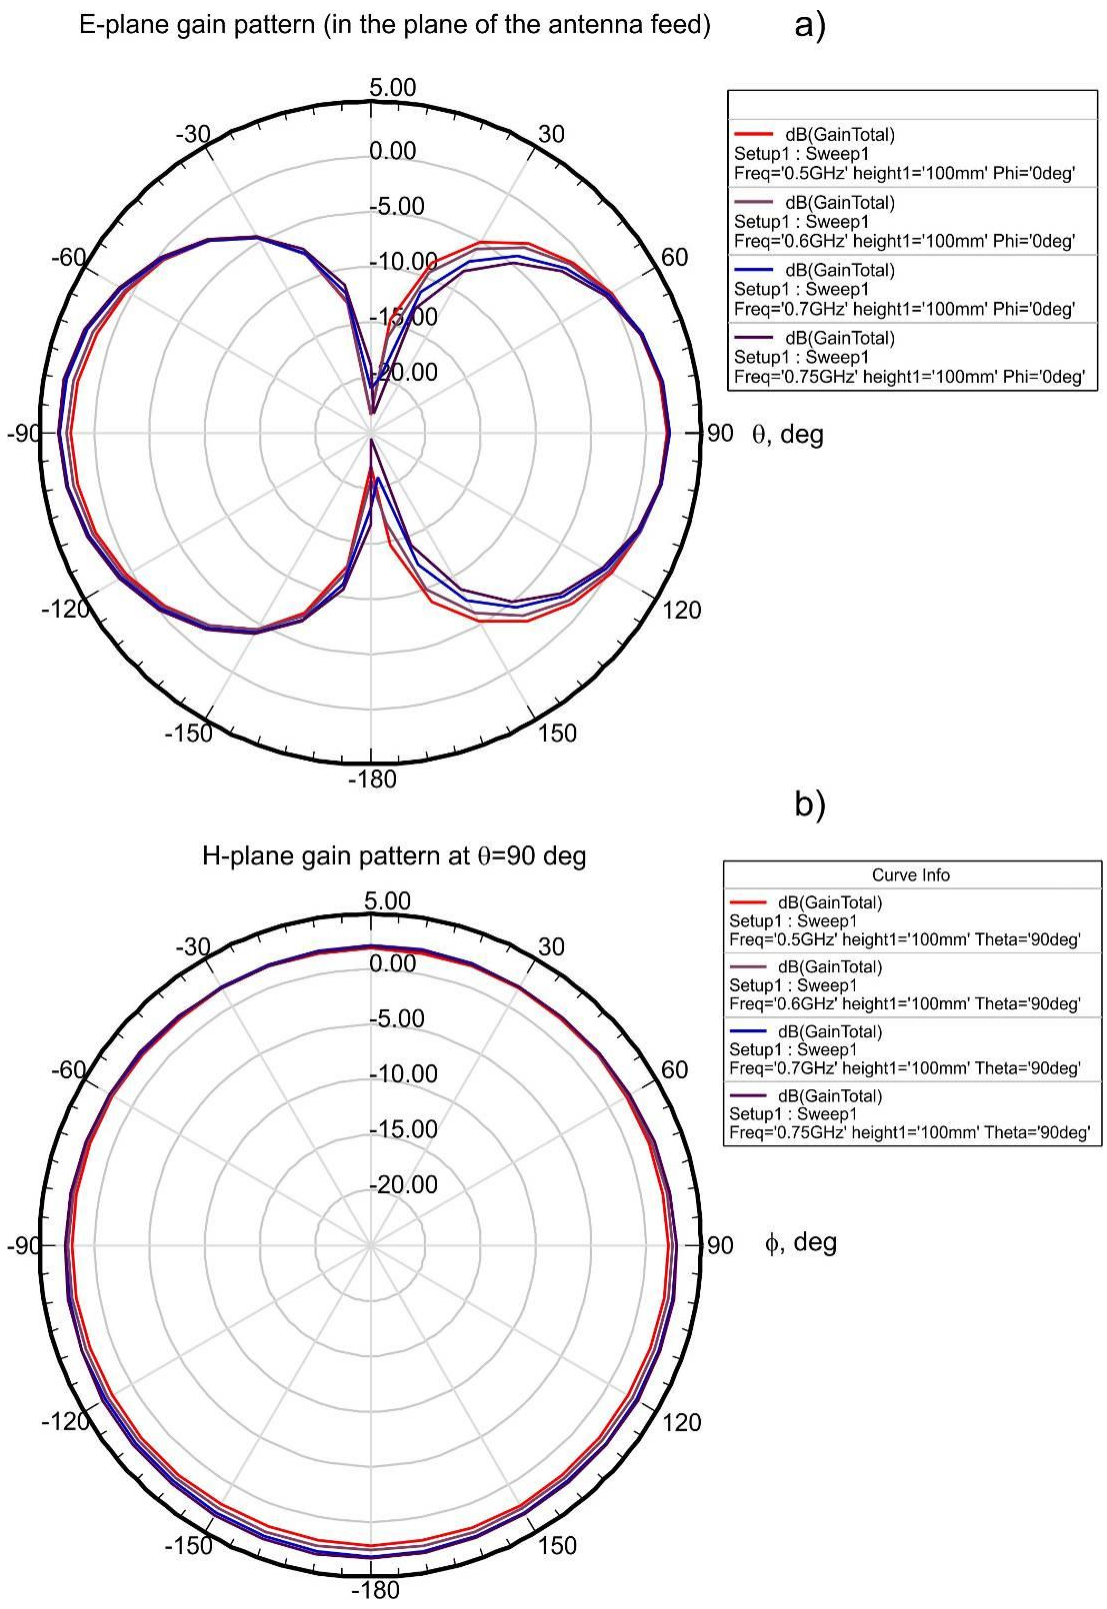

Figure 4.9: Telescopic Antenna Radiation Pattern [16]

Chapter 5

Bayesian Fusion Algorithm Simulation Results

The same set of simulations that was executed as a benchmark for performance of the TART algorithm [5] were executed with the incorporation of the Bayesian Fusion Algorithm. For simulation purposes, we only considered fusion of σART and TART data on a 2D plane. This choice can be thought of as the result of making the assumption of perfect barometric information. Using ideal height information allowed us to see the impact of the fusion of pure RF data, and verify assumptions made in Section 3.2. The results shown here are produced in Matlab, with a simulated PPL waveform being transmitted from the truth location, and received from the reference locations. Reference nodes are shown as white circles, while the truth location is shown as a white square. A black 'X' marks the location where the metric of a particular algorithm is maximized. The color of the image corresponds to the value of the metric function at that point. The lowest values are shown in blue, with the highest values in red. As mentioned above, the absolute value of the metric is not important, and these images are a good visualization of the relative values of the metric.

5.1 Simulated Results

Figure 5.1 shows an ideal simulation with no multipath fading, or noise, and more than the minimum[1] number of antennas needed to solve for the correct location. The antennas are also assumed to be isotropic, and symmetric double solutions are assumed not to be in the search space. In the ideal simulation all three methods correctly identify the location of the transmitter. The area

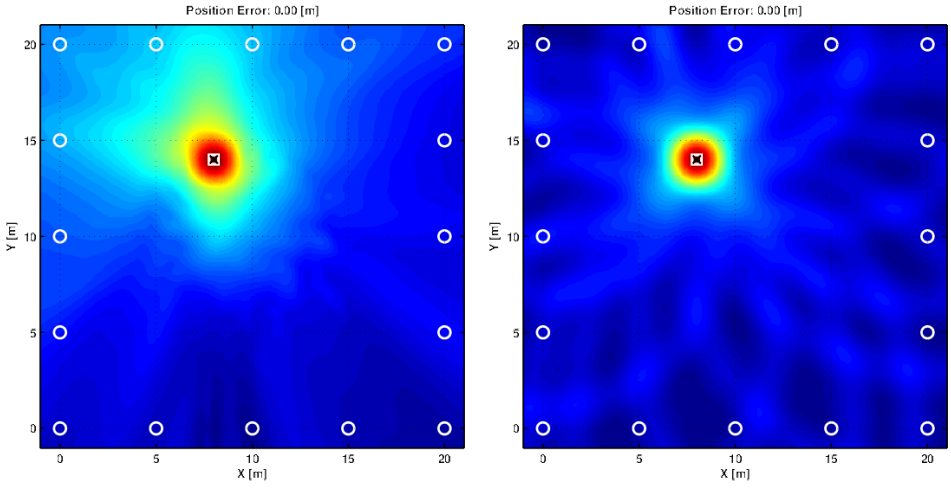

(a) σART [5]

(b) TART [5]

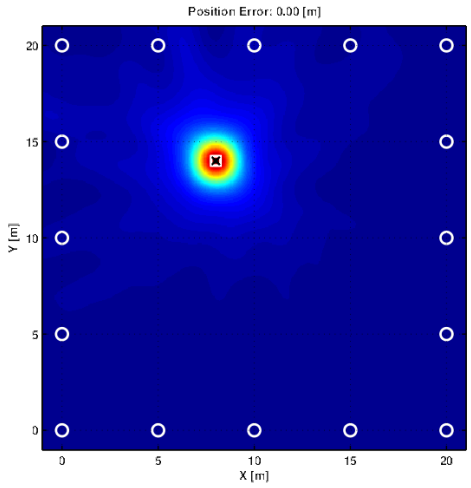

(c) Fusion

Figure 5.1: σART vs. TART vs. Fusion - Ideal Simulation

[1]Two antennas are needed for 2D positioning with the TART algorithm, and three antennas are needed for the σART algorithm.

outside of the peak shows the differences between the different metric functions clearly. The fusion algorithm has the sharpest peak, which makes sense, as the σART and TART algorithms are highly concentrated near the solution region and reinforce each other.

Figure 5.2 shows a simulation in the same configuration as the previous example. Here, we have inserted an ideal reflector, shown by a green triangle. This is essentially another locator that transmits with a time offset proportional to the distance between itself and the true locator. This reflector has no amplitude degradation, and is also assumed to be isotropic. The σART algorithm

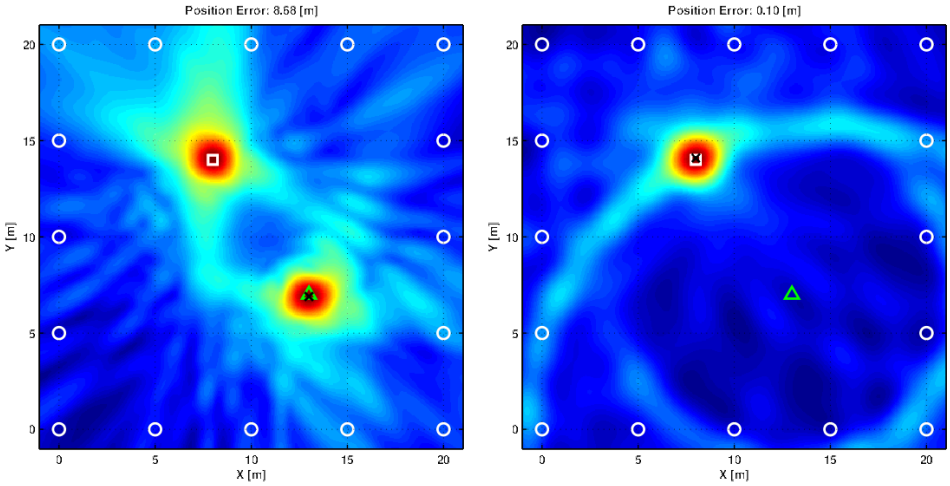

(a) σART [5] (b) TART [5]

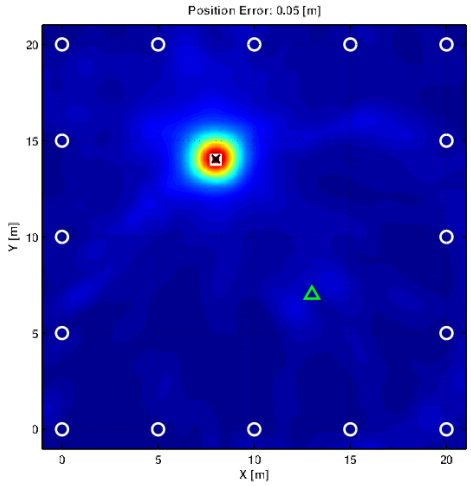

(c) Fusion

Figure 5.2: σART vs. TART vs. Fusion - Single Ideal Reflector

cannot distinguish between this reflector and the real transmitter. The signal from the reflector has an absolute delay compared to the signal from the locator, and the σART algorithm cannot resolve a time delay that is constant at all receive antennas. The resulting error is 8.68 meters. The TART algorithm does a much better job, giving a position estimate that is only 0.1 meters away from the truth location. The TART solution is perturbed by the presence of the reflector because the antennas that have the reflector between them and the transmitter receive both the direct path, and the reflected signal with a relatively small time offset. This creates multiple ridges of maxima in the image, moving the peak slightly. Although the σART metric picked the completely wrong peak, it still had a maximum, at the correct location. This is a situation in which the Fusion algorithm provides an apparently more robust solution, estimating the correct location with an error of just 0.05 meters despite the disturbance introduced by the reflectors.

In Figure 5.3 we add in 6 reflectors in random locations. The σART algorithm completely fails in this situation, with an error of 10.77 meters. There is still a sharp peak at the correct location, but it is smaller in magnitude than the absolute maximum, which is located near a group of reflectors in the lower right hand corner of the image. The TART algorithm returns a solution with 0.4 meters of error, which is quite good, given the high degree of multipath. The peak of the TART metric is rather blurry around the correct location. The Bayesian Fusion Algorithm is able to make use of the fact that the σART input will amplify the true peak, while attenuating the area immediately surrounding it. Furthermore, the TART algorithm is very confident that the locator is not in the region where the σART metric is maximized, which nullifies this false peak, and the fusion result estimates the correct position of the locator within 0.3 meters.

To investigate the effects of a Geometric Dilution of Precision (GDOP) on the different algorithms, we remove the antennas on two sides of the scan grid. In this simulation we assumed an ideal channel, with no reflectors and more than enough antennas to correctly perform location. Figure 5.4 shows the resulting metric functions. As was the case in the first simulation, all three algorithms correctly identify the position of the locator. The σART metric function has a peak which is blurred along a diagonal, while the TART metric function is blurred in a direction that is perpendicular to this diagonal. The Fusion metric function has a circular peak where these two peaks overlap, and the peak is significantly higher than the background metric.

Figure 5.5 shows the type of GDOP that would be experienced if we only had access to one

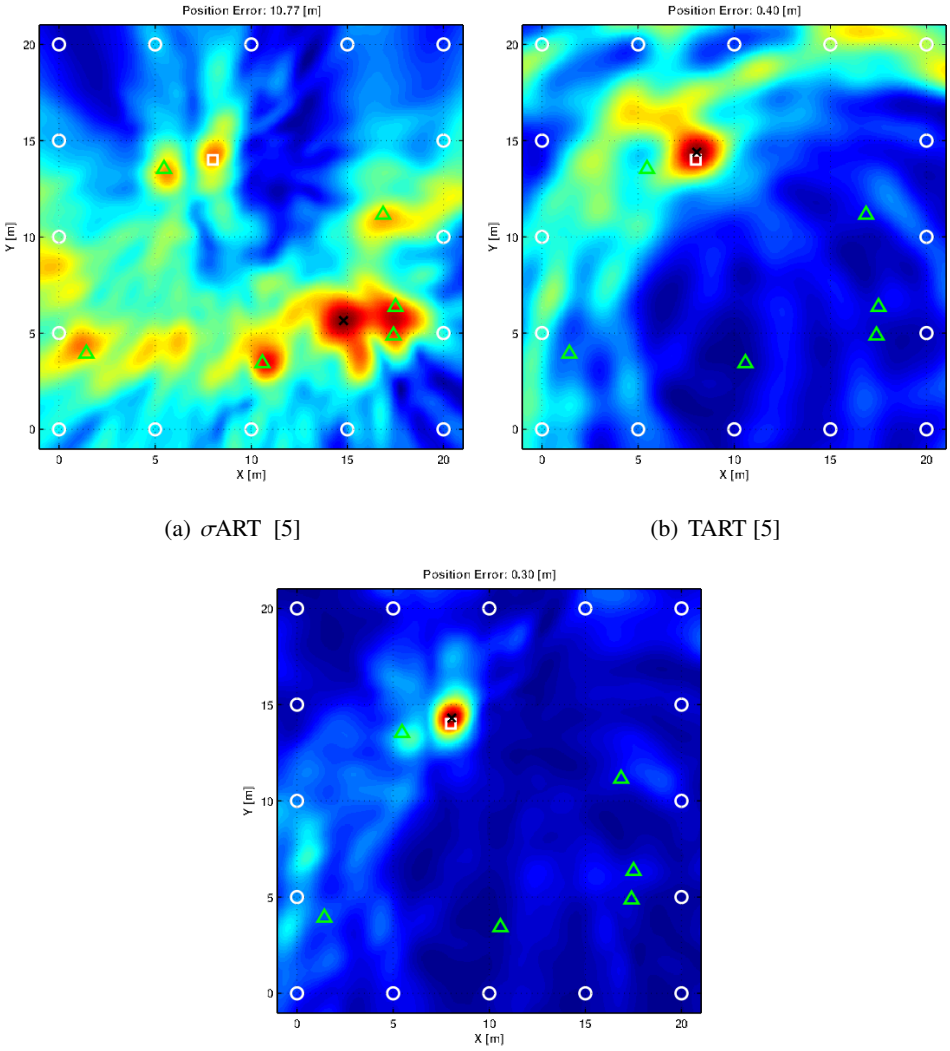

(a) σART [5]

(b) TART [5]

(c) Fusion

Figure 5.3: σART vs. TART vs. Fusion - Randomly Distributed Ideal Reflectors

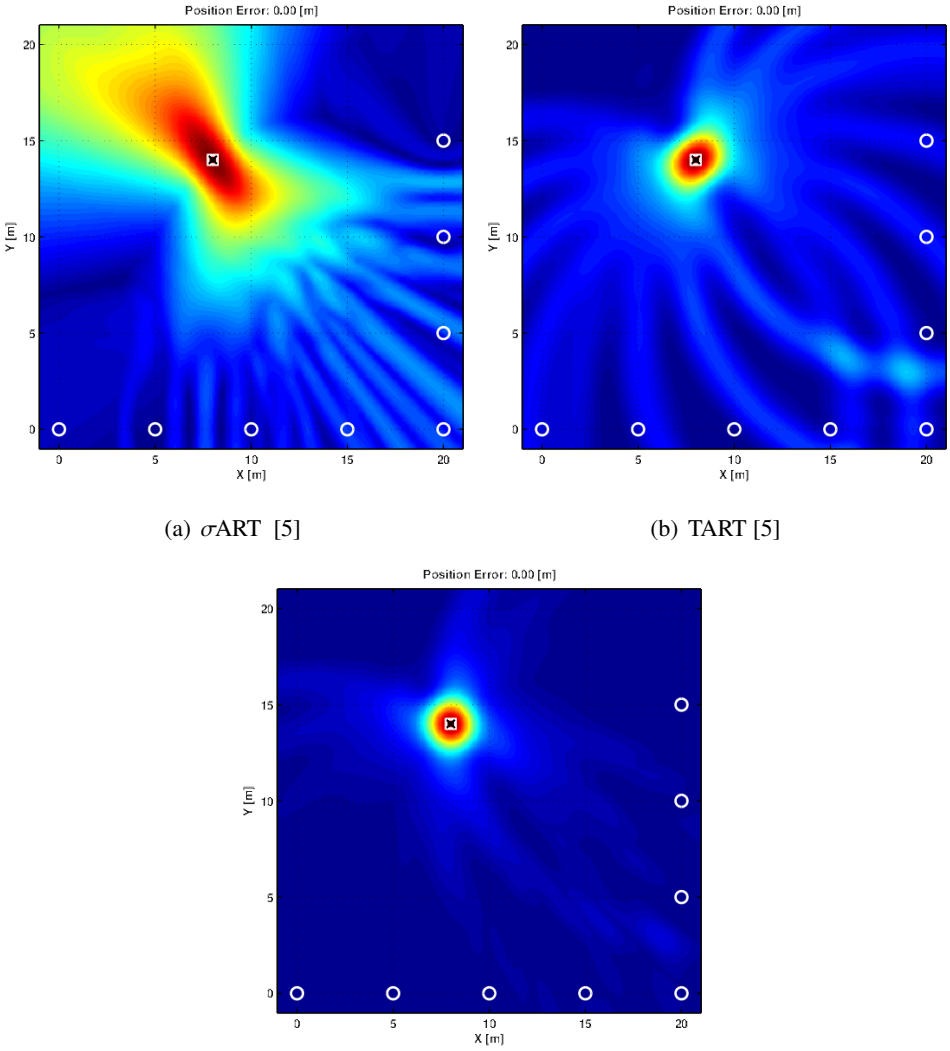

(a) σART [5]

(b) TART [5]

(c) Fusion

Figure 5.4: σART vs. TART vs. Fusion - Geometric Dilution of Precision

side of a building. The metric functions look very different among the three algorithms. The

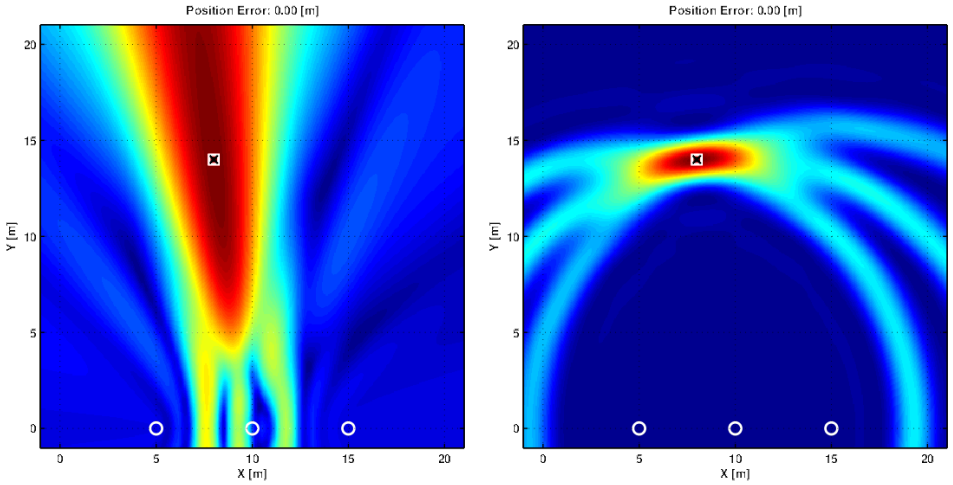

(a) σART [5]

(b) TART [5]

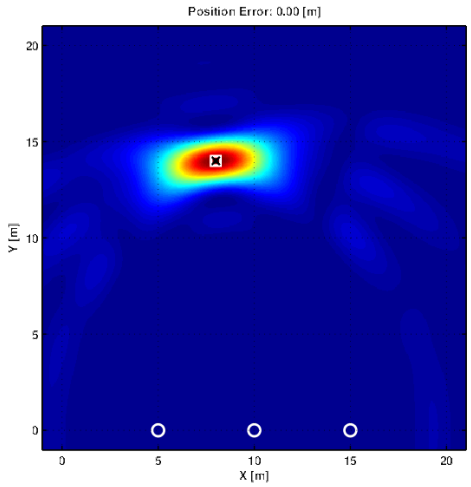

(c) Fusion

Figure 5.5: σART vs. TART vs. Fusion - Severe Geometric Dilution of Precision

σART metric function is blurred in the direction perpendicular to the line formed by the receiving antennas. The TART metric function is blurred along a line parallel to the receiving antennas. Once again, the orthogonal nature of the spreads of the σART and TART metric functions makes the Fusion metric function very focused around the correct solution, which gives us a solution that is more robust in the presence of noise, as seen in Figure 5.6, where the SNR has been decreased from infinity to 18dB. The most likely place for noise to move the solution is along the ridge of maxima

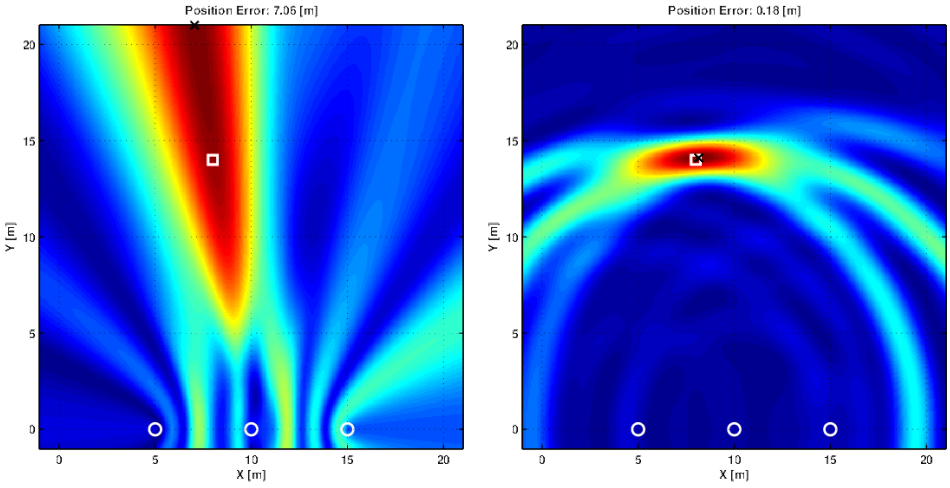

(a) σART [5]

(b) TART [5]

(c) Fusion

Figure 5.6: σART vs. TART vs. Fusion - Severe Geometric Dilution of Precision with 18 dB SNR

in the metric function. Therefore a tighter peak will result in better noise performance. With noisy data it is also possible that one of the other algorithms may "get lucky" and perform better than the fusion, but the fused result would be expected to stay within an acceptable error range in more cases. An example of the "Lucky" phenomenon is shown in Figure 5.7 where the σART algorithm now gives the best result, with an error of 0.05 meters. Here the noise has perturbed the solution to the alternate peak. The TART solution remains the same at 0.10 meters, while the Fusion error grew slightly to 0.10 meters.

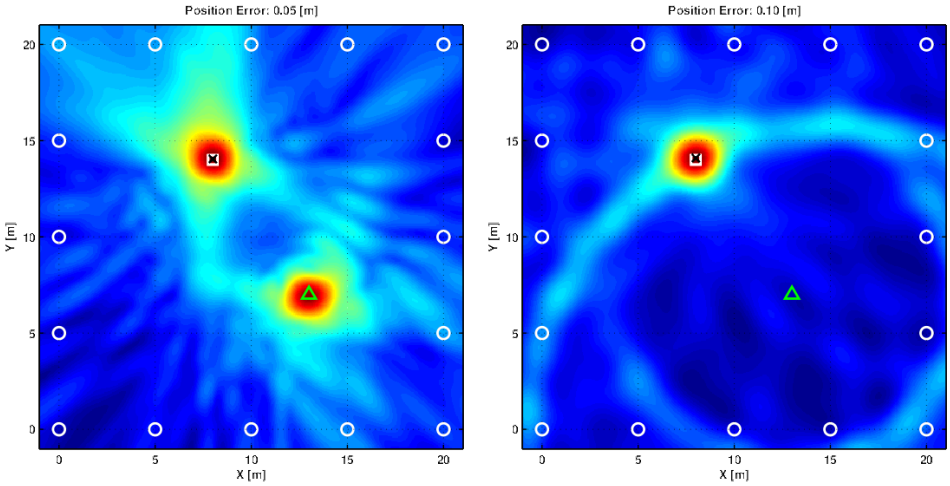

(a) σART [5] (b) TART [5]

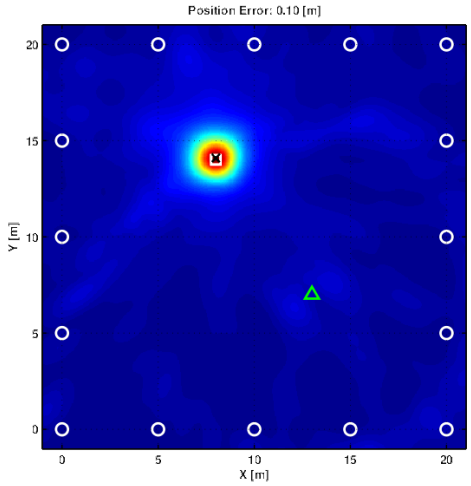

(c) Fusion

Figure 5.7: σART vs. TART vs. Fusion - Single Ideal Reflector with 12 dB SNR

While Figures 5.6 and 5.7 nicely highlight the ability of the Bayesian Fusion Algorithm to cope with the non ideal effects of noise, they are just single cases. Figure 3.2 shows that the result of Monte Carlo testing with 10,000 cases with -6dB SNR, and the results were consistent with those observed in these single test cases. The Bayesian Fusion algorithm produced the correct result more frequently than either σART or TART alone, and in the 10,000 trials it never produced significant outliers; which is not an indicator that the fusion algorithm never produces outliers, it just produces them with a much lower probability than the σART and TART algorithms. The simulations executed here show that the Bayesian Fusion algorithm consistently performs as well or better than σART or TART do individually.

Chapter 6

Experimental Results

This chapter presents the results from two field tests conducted in both residential and commercial settings. The goal of these tests was to demonstrate the rapid deployment of our system, as well as collecting data to be fused in our Bayesian Fusion Algorithm. Unless otherwise stated, all "Fusion" results are full fusions of σART , TART, and barometric data. Two barometric-only tests were also conducted to characterize the performance of our sensor in real-world conditions, such as in a building in which a fire is burning.

When reading this chapter several common types of figures and tables will be used. Test layouts will show a floor plan with labeled squares, which correspond to truth locations for that test. The blue circles in these figures represent the reference antennas. Similar diagrams show the errors at each truth location as a vector which points from the correct location to the algorithm's location estimate. Table 6.6 summarizes the performance of the three algorithms in the two test locations.

The tests in this chapter used a PPL multicarrier waveform with 100 MHz. of bandwidth and 109 carriers transmitting at 10 dBm. The RF data was captured using 64 symbol fusion to boost SNR (functionally equivalent to simple averaging). The bandwidth extrapolation technique referenced in [5] was employed on the σART and TART processing for the Atwater Kent Laboratories testing, and on the σART processing for the Campus Ministry testing.

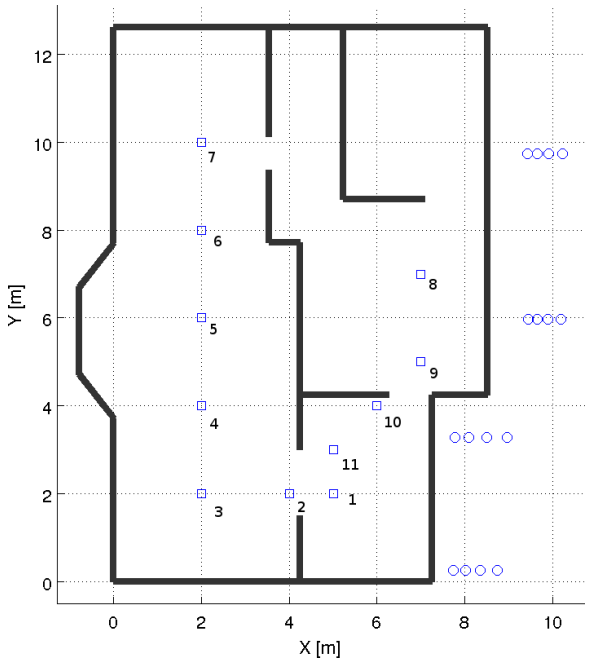

Figure 6.1: Religious Center Truth Locations from 7/31/2009

6.1 Post Processed Fusion from WPI Religious Center

The first test of the rapid deployment scheme was on July 31, 2009, at the WPI Religious Center. Figure 4.1 is a photograph of four ladders with four antennas mounted on each, at this test. At this test site, we used our second generation locator to collect σART data with the transceiver boxes on a common clock so that we could see the effect of this antenna geometry without having to factor in any problems that could have arisen from error in synchronization procedures. Later, that same day, we captured TART data using our fifth transceiver box as a mobile locator, on a separate clock. The locations of the truth points used in these tests is shown in Figure 6.1.

On January 20, 2010, we returned to the religious center to capture the pressure data that was required to process the test data with our Bayesian Fusion algorithm. The original test was not intended to be used for this purpose, so the σART, TART, and pressure data were captured at different times, on different hardware. For the RF tests, it is a good assumption to make that the channel did not change between the σART and TART captures, as no reflectors entered or left the site during the testing on that day. The pressure in January was likely different from that in July, but since we only use differential pressure measurements, which cancels any global effects from weather, these

Location	Truth Height [m]	Measured Height [m]	Error [m]
1	1.19	1.61	0.42
2	1.19	1.61	0.42
3	1.19	1.49	0.30
4	1.19	1.60	0.41
5	1.19	1.69	0.50
6	1.19	1.68	0.49
7	1.19	1.57	0.38
8	1.19	1.72	0.53
9	1.19	1.65	0.46
10	1.19	1.63	0.44
11	1.19	1.62	0.43

Table 6.1: Barometric Height Estimates for Religious Center

captures are likely a good reflection of what we would have seen on July 31, 2009. The heights estimated from the differential pressure measurements are shown in Table 6.1, along with the true heights at the given locations. The σART and TART error vectors from this test are shown in Figure 6.2. The resulting fusion errors are shown in the error vector plot in Figure 6.3, with the corresponding error values listed in Table 6.2. As expected, the errors from the Fusion Algorithm are the smallest in several cases; when they are not the smallest, they are much closer to the minimum error acheived by either σART and TART than they are to the maximum of the σART and TART errors.

6.2 Results from Atwater Kent Laboratories

On May 27, 2010 a test of the PPL system was conducted in the west wing of The Atwater Kent (AK) building on the WPI campus. This test was the first to simultaneously capture σART , TART, and barometric data with the rapid deployment setup. Sixteen patch antennas were affixed to four ladders which were leaned against a wall of the west wing of the building. This wing of the building comprises three stories and a basement. The top half of the basement is above-grade on one side of the building. The structure is brick with a modern steel and dry-wall interior. There are also numerous large pieces of machinery, and no windows, on the western-most wall. A photograph of the reference antennas is shown in Figure 6.6. We collected data on the basement, first, and second

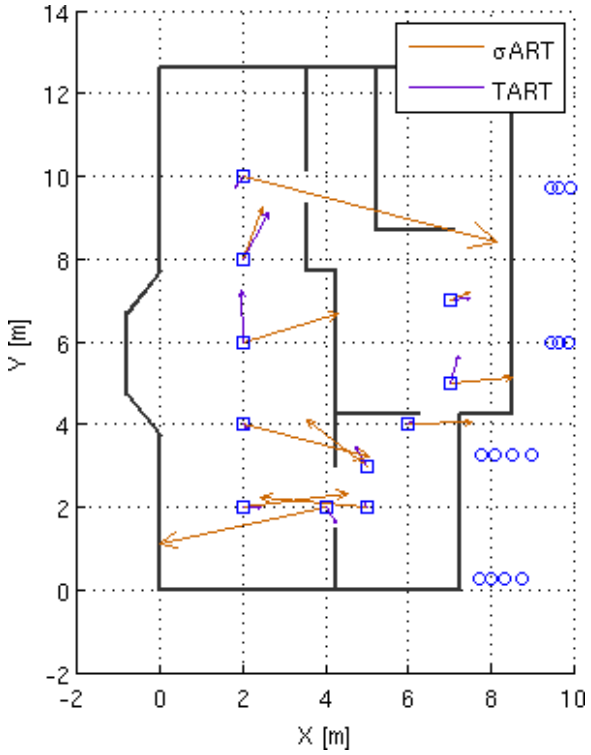

Figure 6.2: σART and TART Error Vectors from 7/31/2009

Location	XY Error [m]			Z Error [m]			XYZ Error [m]		
Method	σART	TART	Fusion	σART	TART	Fusion	σART	TART	Fusion
1	2.57	0.42	0.28	0.16	0.39	0.33	2.58	0.58	0.43
2	3.14	0.11	0.76	0.37	0.77	0.33	3.16	0.78	0.83
3	2.41	1.22	1.03	0.13	0.34	0.33	2.41	1.27	1.08
4	1.34	1.28	1.15	0.33	0.77	0.33	1.38	1.50	1.20
5	6.35	0.31	0.45	1.49	0.09	0.33	6.52	0.33	0.56
6	4.09	0.47	0.56	1.17	0.11	0.33	4.26	0.49	0.65
7	2.57	0.12	0.25	0.39	0.04	0.33	2.60	0.13	0.41
8	1.83	0.55	0.67	0.16	0.51	0.33	1.83	0.75	0.75
9	1.55	0.16	0.13	0.23	0.74	0.33	1.56	0.76	0.35
10	1.51	0.70	0.38	0.17	0.29	0.33	1.51	0.75	0.50
11	0.51	0.47	0.28	0.32	0.45	0.33	0.60	0.65	0.43

Table 6.2: Religious Center Errors for σART , TART, and Fusion Solutions

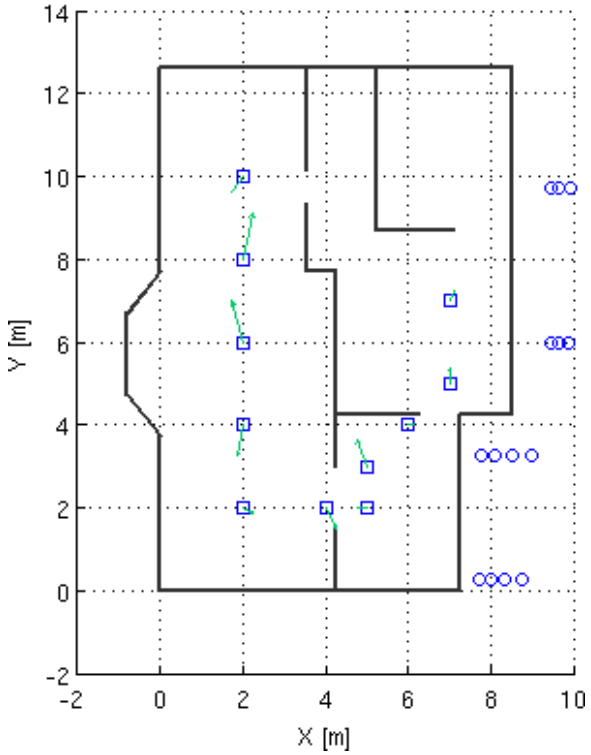

Figure 6.3: Error Vectors from the Post-Processed Religious Center Fusion (7/31/2009)

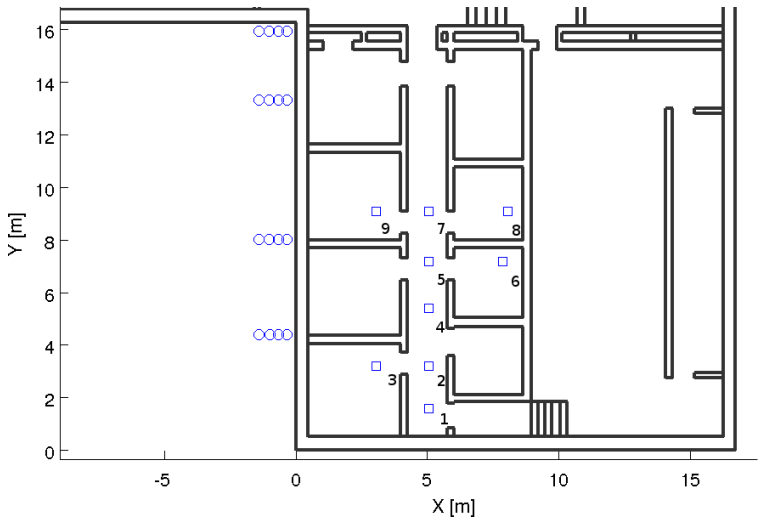

Figure 6.4: Atwater Kent Basement Truth Locations from 5/27/2010

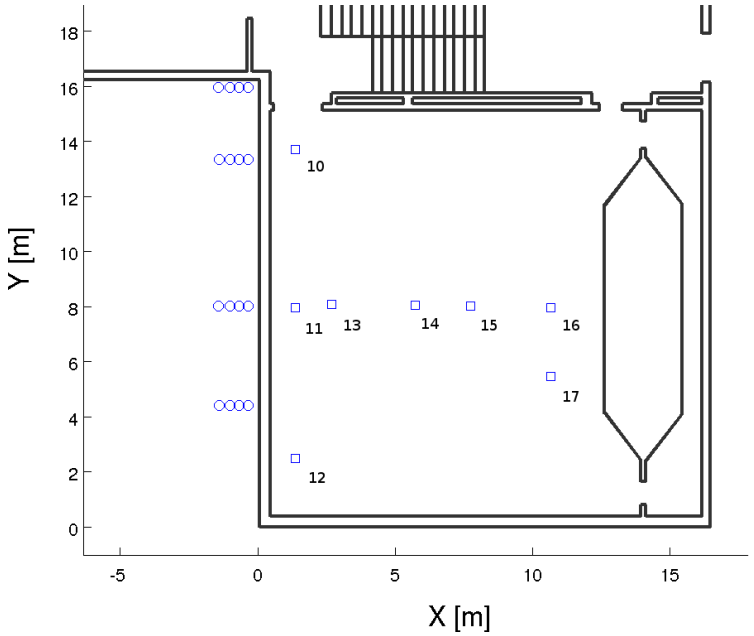

Figure 6.5: Atwater Kent First Floor Truth Locations from 5/27/2010

Figure 6.6: Exterior Photograph of Atwater Kent Laboratories showing the west wing

floors.

The basement of Atwater Kent Laboratories is interesting because the layout is that of a typical office building, but the environment itself is challenging. Challenging RF conditions include being partially below-grade, only having windows on one side of the building, and being surrounded by elevator and HVAC equipment. The layout of the truth points for the basement is shown in Figure 6.4

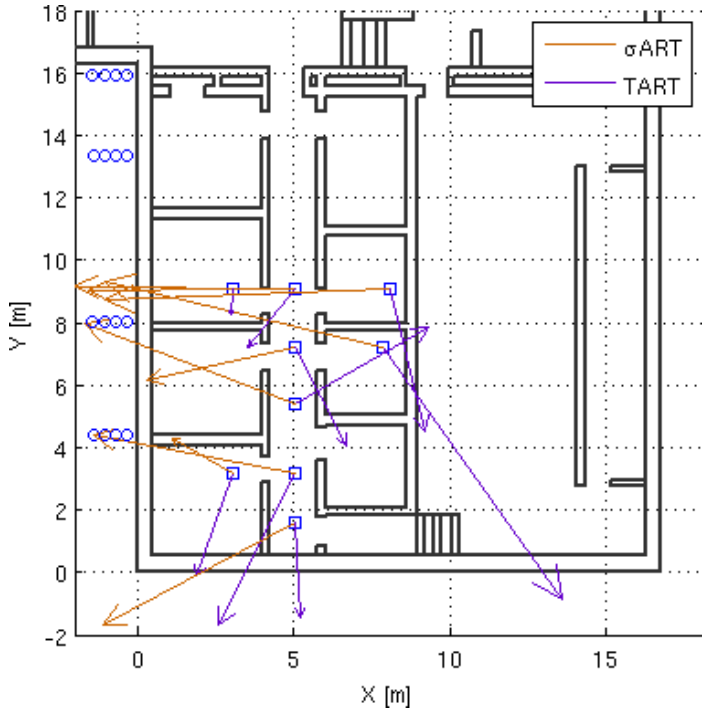

Figure 6.7: σART and TART Error Vectors for AK Basement

The first floor is the most interesting case, because it is a lecture hall with theater seating, which allowed us to place truth locations at many different heights, rather than having uniform, discrete, floor heights. This is the first test setup that truly tests our accuracy in full 3D. The layout of the truth points on the first floor is shown in Figure 6.5.

Although the pressure errors are larger than the required ± 1.5 meters, there seems to be a bias that could be removed based on measurements taken on the basement level. In the case of this test the height estimates that were below the basement elevation rejected, and the scan height was chosen to be near the correct height. In Table 6.5 it is clear that the incorporation of prior information regarding the minimum height on the scan grid resulted in zero error for these cases.

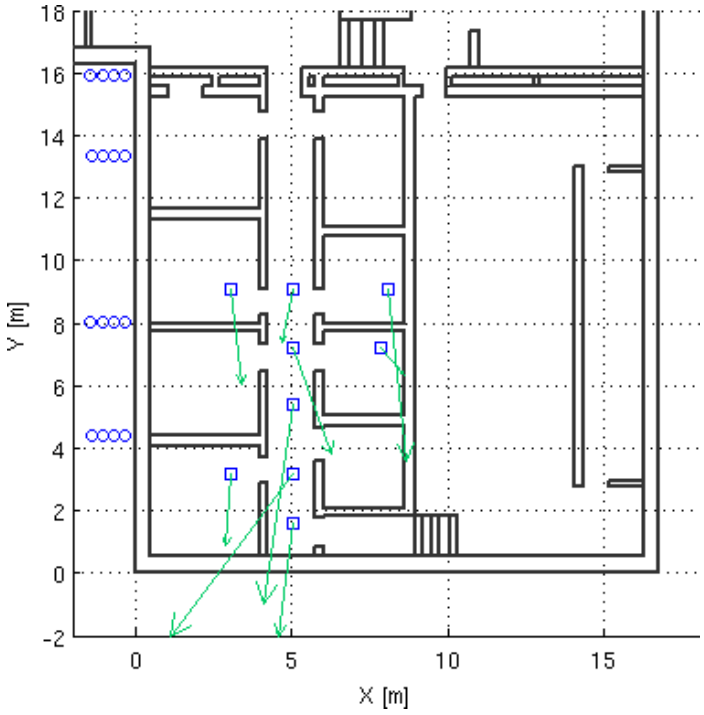

Figure 6.8: Fusion Error Vectors for AK Basement

Location	Truth Height [m]	Measured Height [m]	Error [m]
1	1.19	1.18	-0.01
2	1.19	-0.78	-1.97
3	1.19	0.29	-0.90
4	1.19	-1.01	-2.20
5	1.19	-0.20	-1.39
6	1.19	0.29	-0.90
7	1.19	0.03	-1.22
8	1.19	0.93	-0.26
9	1.19	0.29	-0.90

Table 6.3: Barometric Height Estimates for the Atwater Kent Basement

Location	Truth Height [m]	Measured Height [m]	Error [m]
1	1.19	1.18	-0.01
2	1.19	-0.78	-1.97
3	1.19	0.29	-0.90
4	1.19	-1.01	-2.20
5	1.19	-0.20	-1.39
6	1.19	0.29	-0.90
7	1.19	-0.03	-1.22
8	1.19	0.93	-0.26
9	1.19	0.29	-0.90
10	5.20	1.80	-3.40
11	5.20	2.20	-3.00
12	5.20	2.79	-2.41
13	5.00	2.27	-2.73
14	4.40	2.70	-1.70
15	3.59	3.83	0.24
16	2.96	2.23	-0.73
17	2.96	2.87	-0.09

Table 6.4: Barometric Height Estimates for Atwater Kent 116

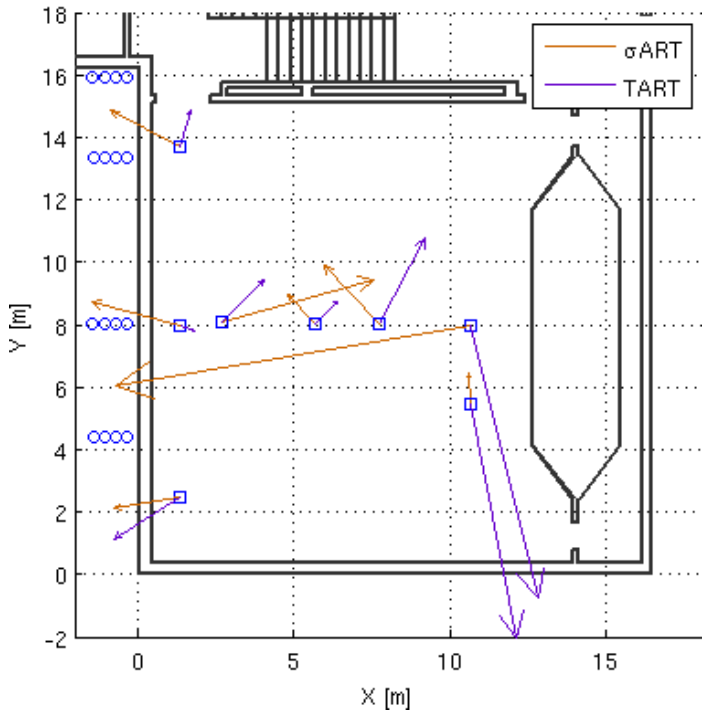

Figure 6.9: σART and TART Error Vectors for AK 116

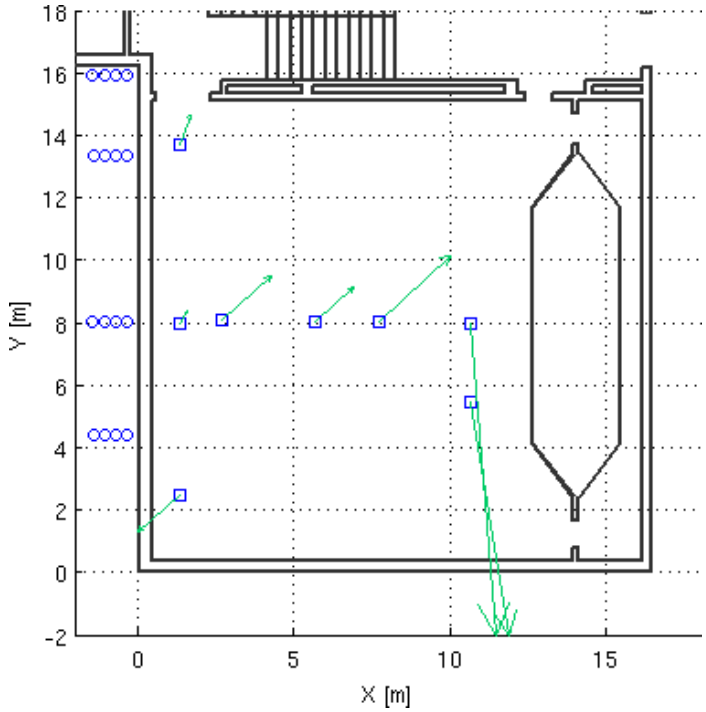

Figure 6.10: Fusion Error Vectors for AK 116

The best example of a successful fusion from this test is at truth point #3, on the first floor of Atwater Kent. The σART and TART metric functions are shown in Figure 6.2, along with the fused result.

6.3 Summary of Experimental Results

Table 6.6 shows the σART only, TART only, and Bayesian Fusion algorithm overall errors from the Campus Religious Center and Atwater Kent Laboratories. In general, the Bayesian Fusion Algorithm improved the performance of the PPL system in the Campus Religious Center, where we expect to have very high SNR and only moderate multipath conditions. In Atwater Kent the Bayesian Fusion Algorithm improved the 2D case, but because of unreliable pressure data, the 3D accuracy was worse. Section 5.1 presented several scenarios in which the Fusion algorithm is not the best solution, however in all cases reviewed, the errors from the Fusion are within the neighborhood of the σART and TART errors. The starred tests (*) are the result of a purely σART and TART fusion computed at the correct height of the locator. These results show that the 2D solution is sensitive to the height estimate used in the selection of the σART and TART scan plane. Contrary to

Location	XY Error [m]			Z Error [m]			XYZ Error [m]		
Method	σART	TART	Fusion	σART	TART	Fusion	σART	TART	Fusion
1	6.96	3.04	3.62	2.75	0.52	0.00	7.48	3.08	3.62
2	6.58	5.46	6.53	1.36	0.77	0.00	6.72	5.51	6.53
3	2.24	3.51	2.32	0.77	0.00	0.00	2.37	3.51	2.32
4	7.24	4.90	6.42	0.12	0.00	0.00	7.24	4.90	6.42
5	4.87	3.52	3.61	0.52	0.52	0.00	4.90	3.56	3.61
6	9.17	9.87	1.14	2.93	6.18	0.00	9.63	11.65	1.14
7	7.06	2.42	1.73	0.00	5.12	0.00	7.06	5.66	1.73
8	9.11	4.71	5.53	3.02	0.00	0.00	9.60	4.71	5.53
9	4.60	0.81	3.07	0.82	6.95	0.00	4.67	7.00	3.07
10	2.47	1.21	1.67	0.06	0.58	3.49	2.48	1.34	3.87
11	2.93	0.50	0.31	0.81	0.48	2.46	3.03	0.69	2.48
12	2.12	2.52	0.74	0.38	0.20	2.46	2.16	2.53	2.57
13	5.07	1.95	1.57	1.41	0.66	2.78	5.27	2.04	3.20
14	1.26	0.98	1.01	1.42	1.92	1.66	1.90	2.16	1.95
15	2.58	3.12	3.08	1.73	1.88	0.18	3.11	3.65	3.08
16	11.48	8.97	9.87	2.90	4.15	0.75	11.84	9.89	9.90
17	0.97	7.61	7.57	1.77	0.29	0.03	2.02	7.61	7.57

Table 6.5: Atwater Kent Errors for σART , TART, and Fusion Solutions

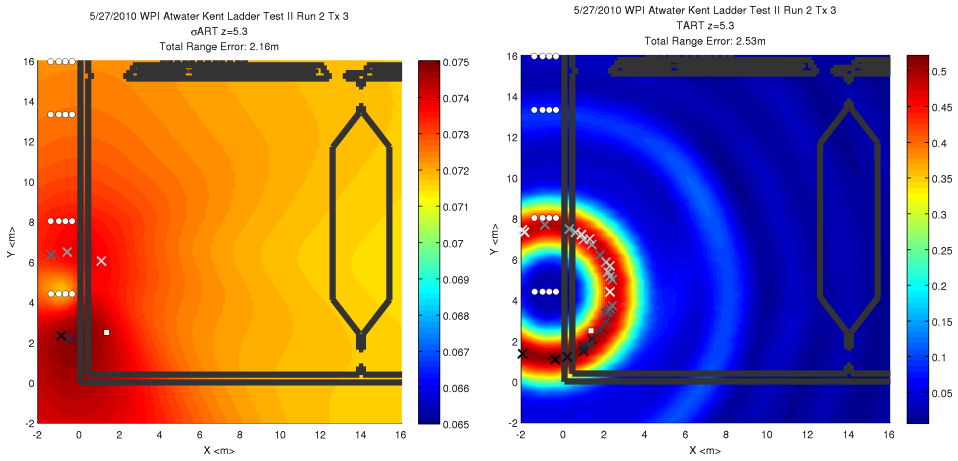

(a) σART [5] (b) TART [5]

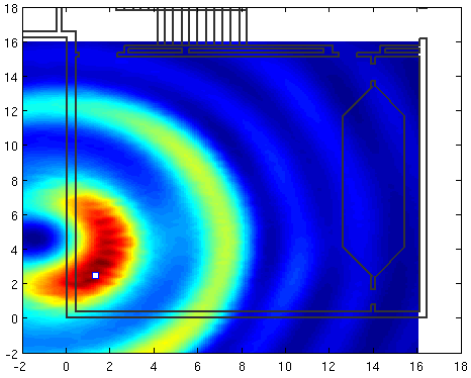

(c) Fusion

Figure 6.11: σART vs. TART vs. Fusion - AK Ladders Tx 3

Test	Method	XY Mean Abs. Error [m]	Z Mean Abs. Error [m]	XYZ Mean Abs. Error [m]
Religious Center	σART	2.53	0.51	2.58
	TART	0.53	0.49	0.72
	Fusion	0.54	0.36	0.65
Religious Center*	σART	1.75	0.00	1.75
	TART	0.49	0.00	0.49
	Fusion	0.57	0.00	0.57
Atwater Kent	σART	3.61	1.67	3.98
	TART	3.36	1.64	3.74
	Fusion	3.23	2.88	4.33
Atwater Kent*	σART	7.91	0.00	7.91
	TART	4.51	0.00	4.51
	Fusion	4.72	0.00	4.72

Table 6.6: Summary of Errors from Methods Presented

naive expectations the lowest 2D errors were not always found when scanning at the correct height. In the Atwater Kent case the TART error alone increased (by about a meter) when scanning at the correct height rather than the barometric estimate of height. This is likely due to the dispersion of the likelihood function that we observe in the TART metric function when some of the multicarrier signals are delayed by significant obstructions, such as brick walls. Since this delay is not nearly as pronounced in the Religious Center, scanning at the correct height should increase the accuracy of TART. The increase in accuracy here would only be modest because the barometric height estimates were relatively close to the correct value, and was not apparent in the results in which a slight increase in total error was actually seen.

Chapter 7

Conclusion

The Bayesian Fusion Algorithm was developed to recover accuracy that was lost when we began deploying the PPL system with realistic restrictions imposed on the placement of reference antennas. The distinct error characteristics of TDOA-like and TOA-like algorithms initially led us to consider Bayesian fusion [13] for overall accuracy improvement. To perform the same type of fusion seen in [13] and [10] we had to show that the probability distributions of σART and TART solutions were statistically independent. The accuracy of the Bayesian Fusion algorithm was then tested with both simulations and with field experiments. In addition to improving accuracy, the Bayesian Fusion Algorithm was expected to increase the robustness of positioning solutions when perturbed by noise, which was verified in simulation, and was supported by experimental data.

A shortcoming of the Bayesian Fusion Algorithm is that there is no known way to assign a confidence metric to the underlying data sources. All of the barometric sensors that we have tested have been fraught with errors and calibration issues of every kind, and neither the σART nor the TART algorithms provide us with a measure of confidence in their solutions. The addition of this type of confidence information would likely increase the accuracy and robustness of the Bayesian Fusion Algorithm. A major strength of the Bayesian Fusion Algorithm is that it does provide us with a measure of it's own confidence based on the correlation between the underlying data sources. This means that the incorporation of additional data sources with the current algorithm could benefit from this type of information. The system being developed within the inertial navigation research being done in parallel with this thesis [22] is an example of another data source that could be included

in such a fusion. There is also a great deal of promise in using this algorithm for tracking moving locators, as opposed to single position estimates alone to improve overall accuracy in conjunction with a Kalman filter, as information from the past can benefit the solution in the present time step.

7.1 Contributions of this Thesis

7.1.1 Utility of the σART Algorithm

The TART algorithm was initially seen as a direct replacement for the σART algorithm, a next step in improving the accuracy of the PPL system. Simulations and field tests in this and other works [5] have consistently shown the TART algorithm to perform well in situations where the performance of the σART algorithm is compromised. The question remains whether the inverse statement is true: does the σART algorithm perform well in situations where the performance of the TART algorithm is compromised? Furthermore, in situations where both algorithms are compromised, can they each provide useful information for a fused result?

There is reason to believe that the σART algorithm should outperform the TART algorithm in certain situations. Consider a situation in which we have reference antennas on the outside of a concrete structure, and a locator within the walls of this structure. The fundamentally TOA-like TART algorithm cannot compensate for the delay caused by the lower speed of light within the concrete. The σART algorithm, however, is not affected by this delay, as it is constant for all of the reference antennas, and therefore produces no *difference* in the time of arrival. The simulations presented in Chapter 5 also show several instances where σART outperforms TART. Experimental results show that in cases where both algorithms perform poorly, the data provided by σART and TART are distinct and provide us with useful information. This observation led to the idea that TART complements σART, rather than superseding it.

7.1.2 Statistical Independence of σART and TART Errors

Once it was established that σART and TART could both provide us with useful information which could lead to a better, federated solution we focused on Bayesian fusion methods. We believed that this type of fusion would make the best use of the available data, but in order to use Bayesian fusion methods from past work [13, 10] these data had to be shown not fully dependent.

Since the error characteristics were known to be Gaussian [5], it was sufficient to show that the errors were not fully correlated, which seemed likely, based on observations of σART and TART solutions from the same RF data, which showed the error distributions of σART and TART (which are proportional to the metric functions) to be orthogonally skewed. While such skewness is neither sufficient nor necessary, it can be an informal indicator for estimators of the same quantity.

Monte Carlo simulations were conducted in order to calculate the correlation between the σART and TART errors, and it was found that they were not highly correlated. The highest correlation values were found when the SNR was lowest, and the errors were smallest. These cases are where our assumptions of independence are boldest, however, since the individual errors are very small the negative impact of being incorrect in our assumption is minimal. In cases where the correlation is small, our assumptions clearly hold, the errors are larger and hence gains are possible. When the correlation drops to near zero, our assumptions are most valid, but we are likely operating in a region where the errors may be so large as to possibly invalidate other aspects of our statistical model (Gaussianity for example) hence let the performance gain again be available. Thus the advantages to be had from fusion would ultimately depend on upon the behavior of real data from physical environments.

7.1.3 Gains in Accuracy from the Bayesian Fusion Algorithm

Ultimately, the goal of the Bayesian Fusion Algorithm is to improve the accuracy of the PPL system. The gains in accuracy that were observed in simulation were consistent with our expectation that positioning accuracy would increase. The performance in the presence of noise, reflectors, and poor antenna coverage were generally better than the solutions from the σART or TART algorithms on their own. In no cases did the fused result look significantly worse than σART or TART, and it was never the worst of the three.

In the Atwater Kent field tests the Bayesian Fusion Algorithm performed better (in the sense of mean absolute error) than either the σART or TART algorithms. It is also noteworthy that none of the algorithms did particularly well in this very challenging scenario. The fused result did, however, narrow the search to a single office and its nearest neighbors. The barometric data from this test was extremely bad, but additional analysis in which we provided perfect height information showed that the 2D error was actually increased, which speaks to the sensitivity of positioning solutions to scan

height, but in a non intuitive manner.

In the Religious Center, the performance of the Bayesian Fusion Algorithm (using the measure of mean absolute error) was approximately equal to that of the TART algorithm. Unlike Atwater Kent, the religious center does not have thick masonry walls which adversely affect the TOA calculations needed for TART synchronization. This is an example of a situation where σART is totally compromised, but TART is able to perform very well. It is important to note that although the σART solutions were extremely bad, they did not cause the fused result to be bad, which lends support to the claimed robustness of the overall fused solution.

7.1.4 Gains in Robustness from the Bayesian Fusion Algorithm

As seen in the field, large errors from a single data source are not sufficient to corrupt the overall fused solution. The Bayesian Fusion Algorithm is fundamentally robust to error from a single data source because it is evaluating the agreement between not fully dependent measurements on the location of our locator. Monte Carlo simulations also show the Bayesian Fusion Algorithm to be much more robust than σART and TART, with only four position estimates in 10,000 trials that had errors larger than the scan resolution (5 cm), where σART and TART had thousands of incorrect position estimates[1].

7.2 Future Work

In order to further increase the accuracy of the Bayesian Fusion Algorithm, we should shift our focus from evaluating absolute accuracy at static locations to precisely tracking the *path* along which a locator is moving. Given a reasonable estimate of an initial position, and knowledge of how far the user can travel in a time step, it is possible to use a much more informative prior probability distribution to restrict the search space to small area of high likelihood. As time progresses, this prior distribution will grow with the uncertainty of the location, but the growth will be bounded by the current uninformative prior. We can further restrict the growth of this uncertainty by developing a confidence metric for the Bayesian Fusion Algorithm based on the amount of agreement between

[1] These estimates were incorrect, but the degree by which they were wrong was generally small in this ideal simulation with noise.

the σART and TART algorithms.

This type of tracking is also very well suited to integration with inertial navigation data, which has very well known uncertainty characteristics [22]. The addition of inertial navigation data to the Bayesian Fusion Algorithm would further increase the accuracy and robustness of the overall solution, and given a good measure of confidence for the RF results the drift of the inertial system could be kept in check. Since the Bayesian Fusion algorithm (with or without inertial data) can provide confidence information about itself, a Kalman [19] filter could be used to further improve real time path tracking through a building.

Bibliography

[1] Amendolare, V. Synchronization in an Indoor Precision Location System. Master's thesis, Worcester Polytechnic Institute, 2007.

[2] Amendolare, V. Ppl system synchronization. Tech. rep., Worcester Polytechnic Institute, Jan. 2009.

[3] Amendolare, V., Cyganski, D., and Duckworth, R. J. WPI Precision Personnel Location System: Synchronization of Wireless Transceiver Units. In *JSDE/ION Joint Navigation Conference* (June 2009).

[4] Amendolare, V., Cyganski, D., Duckworth, R. J., Makarov, S., Coyne, J., Daempfling, H., and Woodacre, B. WPI Precision Personnel Locator System: Inertial Navigation Supplementation. In *IEEE/ION Position Location and Navigation Symposium (PLANS) Conference* (May 2008).

[5] Amendolare, V. T. *Transactional Array Reconcilliation Tomography for Precision Indoor Location*. PhD thesis, Worcester Polytechnic Institute, 2010.

[6] Anderson, J. R. Abandoned cold storage warehouse multi-firefighter fatality fire. Tech. rep., US Fire Administration, Dec. 1999.

[7] Bard, J. D., Ham, F. M., and Jones, W. L. An Algebraic Solution to the Time Difference of Arrival Equations. In *Southeastcon '96. 'Bringing Together Education, Science and Technology'., Proceedings of the IEEE* (April 1996), pp. 313–319.

[8] Bretthorst, G. L. *Bayesian Spectrum Analysis and Parameter Estimation*. Springer-Verlag, New York, 1988.

[9] Campbell, M. C. Design of a Mobile Transceiver for Precision Indoor Location. Master's thesis, Worcester Polytechnic Institute, 2010.

[10] Cavanaugh, A., Lowe, M., Cyganski, D., and Duckworth, R. J. WPI Precision Personnel Location System: Rapid Deployment Antenna System and Supporting Algorithms for 3D Precision Location. In *ION International Technical Meeting* (Jan. 2010).

[11] Coppock, D. S. The bayesian debate. *Information Management Online* (2003).

[12] Coyne, J. FPGA-Based Co-processor for Singular Value Array Reconciliation Tomography. Master's thesis, Worcester Polytechnic Institute, 2007.

[13] Cyganski, D., Duckworth, R., and Cavanaugh, A. Final report: Microlight indoor positioning performance evaluation sponsered by raytheon company. Tech. rep., Worcester Polytechnic Institute, Jan. 2010.

[14] Dresher, R. Wearable forehead pulse oximetry: Minimization of motion and pressure artifacts. Master's thesis, Worcester Polytechnic Institute, 2006.

[15] Duckworth, J., Cyganski, D., Makarov, S., Michalson, W., Orr, J., Amendolare, V., Coyne, J., Daempfling, H., Hubelbank, D., Parikh, H., and Woodacre, B. WPI precision personnel locator system: Evaluation by first responders. In *Proceedings of ION GNSS* (Fort Worth, Texas, September 2007).

[16] Iyer, V., Cavanaugh, A., Makarov, S., Duckworth, R. J., and Cyganski, D. Self-supporting coaxial antenna with an integrated balun and a linear array thereof.

[17] Jackson, L. B. *Signals, Systems, and Transforms*. Addison Wesely, 1988.

[18] Jaynes, E. T. Prior Probabilities. *IEEE Transactions on Systems Science and Cybernetics 4*, 3 (1968), 227–241.

[19] Kalman, R. E. A new approach to linear filtering and prediction problems. *Transaction of the ASME-Journal of Basic Engineering 82(Series D)* (1960).

[20] Kay, S. M. *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice-Hall, Upper Saddle River, NJ, 1993.

[21] Keshner, M. S. $\frac{1}{f}Noise$. In *Proceedings of the IEEE* (Mar. 1982).

[22] Lowe, M. Inertial Kalman Filtering with Applications in Indoor Localization. Master's thesis, Worcester Polytechnic Institute, 2010.

[23] Nangle, R. Worst building fire disaster in 21 years. *Worcester Telegram and Gazette* (Dec. 1999).

[24] Sivia, D. *Data Analysis*. Oxford University Press, 2006.

[25] Woodacre, B. W. *Geometric Autoconfiguration for Precision Personnel Location*. PhD thesis, Worcester Polytechnic Institute, 2009.